NASA Contractor Report 198172

🌐 ICASE

SEMIANNUAL REPORT

October 1, 1994 through March 31, 1995

Contract NAS1-19480
June 1995

Institute for Computer Applications in Science and Engineering
NASA Langley Research Center
Hampton, VA 23681-0001

USRA   Operated by Universities Space Research Association

CONTENTS

INTRODUCTION

The Institute for Computer Applications in Science and Engineering (ICASE) is operated at the Langley Research Center (LaRC) of NASA by the Universities Space Research Association (USRA) under a contract with the Center. USRA is a nonprofit consortium of major U. S. colleges and universities.

The Institute conducts unclassified basic research in applied mathematics, numerical analysis and algorithm development, fluid mechanics, and computer science in order to extend and improve problem-solving capabilities in science and engineering, particularly in the areas of aeronautics and space research.

ICASE has a small permanent staff. Research is conducted primarily by visiting scientists from universities and industry who have resident appointments for limited periods of time as well as by visiting and resident consultants. Members of NASA's research staff may also be residents at ICASE for limited periods.

The major categories of the current ICASE research program are:

- Applied and numerical mathematics, including multidisciplinary design optimization;

- Theoretical and computational research in fluid mechanics in selected areas of interest to LaRC, such as transition, turbulence, combustion, and acoustics;

- Applied computer science: system software, systems engineering, and parallel algorithms.

ICASE reports are considered to be primarily preprints of manuscripts that have been submitted to appropriate research journals or that are to appear in conference proceedings. A list of these reports for the period October 1, 1994 through March 31, 1995 is given in the Reports and Abstracts section which follows a brief description of the research in progress.

RESEARCH IN PROGRESS

APPLIED AND NUMERICAL MATHEMATICS

SAUL ABARBANEL

Bounded-error methods for long time integration

Work has continued on improving the cell-Reynolds-number restriction which results from the interaction between the advection and diffusion terms in the Navier-Stokes equations.

For the 1-D linear advection-diffusion problem, using second order central differencing, we were able to pose boundary conditions (using the SAT methodology) so that for high Reynolds numbers ($R > 10^4$, say) there are hardly any spurious oscillations in the numerical solution. This situation improves with increasing Reynolds numbers. Using standard boundary conditions, the situation is reversed and the spurious oscillations only increase with R.

The current effort is to generalize the procedure to the multi-dimensional case. This work was carried out with Adi Ditkowski, a graduate student.

EYAL ARIAN

Smoothing analysis for optimization problems

The objective of this project is to develop efficient numerical methods to solve large scale optimization problems which are governed by partial differential equations (PDE's).

The approach is to use multigrid methods where the optimization problem is represented on a set of grids with increasing level of refinement. In this work the multigrid one-shot optimization method was extended to the infinite dimensional design space and applied numerically to optimal shape design problems which are governed by an elliptic PDE. The convergence of the algorithm is influenced by the smoothing properties of the minimization process. This led to the development of a simple Fourier analysis for an optimization problem which computes the symbol of the mapping between the error in the design variables and the sensitivity gradients, (residuals of the design equation) in the high-frequency regime. It shows that the condition number of the Hessian for a given problem depends on the problem set-up, i.e. cost-function definition and the discretization scheme which is used to solve the problem. The analysis can help to construct preconditioners for problems in which the gradients are not sensitive to high-frequency errors in the design variables. It can also help find a preferable discretization scheme and cost function definition to provide a smoothing minimization process for multigrid applications or possibly a well conditioned optimization problem. This work was done in collaboration with Shlomo Ta'asan.

Future plans include the extension of this work to non-elliptic systems.

Fluid-structure optimization problems

The interaction between fluid and structure is especially challenging from the optimization point of view. The problem consists of solving the flow equations around a flexible wing which deflects

1

as a result of the aerodynamic forces. The computation of the aeroelastic steady-state of this system is in itself a difficult problem. The optimization problem seeks wing shape and structure characteristics which are optimal in some sense. The objective of this work is to develop efficient numerical techniques to cope with multidisciplinary optimization problems of this type.

A series of optimization problems which involve structure and fluids is being defined and new numerical techniques for their solution are under development. Preliminary work has been done on a simple model problem involving flow over a panel in a duct.

Future plans include an optimization problem involving flow over a flexible plate in a set-up which models flow over a 3D wing.

JOHN A. BURNS

Feedback control for nonlinear systems

The development of feedback controllers for fluid flows requires that considerable progress be made in the construction of practical low order dynamic compensators. Since the basic models for such systems are nonlinear partial differential equations, one must find low order lumped parameter models that approximate the dynamics of these partial differential equations in some average sense. One approach is to combine linear feedback with a nonlinear observer. The objective of this work is to develop a mathematical framework and computational methods for attacking this problem and to apply the results to problems in fluid flow and structural control.

The approach we consider makes no prior assumptions regarding the form of the controls/actuators. In particular, we assume that the input operator is as general as possible and consider the problem of constructing kernels for integral representations of feedback control laws obtained by solving LQR and MinMax control problems. We have shown how to use these kernels to shape, design and locate sensors. Moreover, in recent work with Belinda King at Oregon State we have been able to use these representations to guide the development of local "reduced basis" nonlinear observers for hyperbolic systems. In order to take full advantage of this methodology one must establish the existence of integral representations for feedback control laws and develop practical computational methods to "approximate" these kernels. We have established the existence of these integral representations and the smoothness of the corresponding integral kernels for parabolic and certain damped hyperbolic systems. In addition, we have applied these results to weakly nonlinear problems with considerable success. Recently, we have conducted a numerical study of the controlled Navier-Stokes equations using this method. The early results for the Navier-Stokes equations are very encouraging. In particular, we could clearly identify regions in the spatial domain where local nonlinear lumped models are sufficient for state estimation. This analysis allows us to develop computational schemes for boundary control problems and in other distributed control problems associated with flow control.

These results will be extended to robust control designs and numerical approximations in order to study the effects of replacing fully distributed observers with local low order lumped observers. We plan to use reduced basis methods and investigate the POD methods developed by Berkooz and others at Cornell. The goal is to apply this method to design practical nonlinear controllers for nonlinear distributed parameter systems.

JAMES GEER

Periodic basis function with built-in discontinuities

Series expansions of functions, such as Fourier series and perturbations series are often useful in developing numerical or semi-numerical semi-analytical algorithms for the solution of differential equations. However, when only a partial sum of such a series is used, some "undesirable" effects such as the Gibbs phenomena may be present, or the partial sum may have "difficulty" approximating certain features of the solution such as boundary layers, internal layers, or various discontinuities.

Several new classes of periodic basis functions which have "built-in" singularities have been defined and some of their fundamental properties have been studied. In particular, it has been shown that for a certain class of functions $f(x)$, the combination of a finite sum of these functions and a finite Fourier series leads to the sequence of approximations which converges exponentially to f, even though f may have a finite number of discontinuities. In particular, these approximations eliminate the Gibbs phenomena.

These new (basis) functions will be studied within the context of developing some "new" spectral methods for systems of ODE's and PDE's. In particular, they will be applied to several model problems which either have discontinuities in the initial data and/or develop discontinuities (or "near-discontinuities") as time increases.

Hybrid Pade-Galerkin techniques

The practical usefulness of perturbation solutions to physical problems is often limited, either because the expansion parameter has to be restricted to "very small" values, or because the solutions are not uniformly valid. In addition, the radius of convergence of perturbation series is sometimes determined by mathematical singularities which are not relevant to the physical problem of interest; for example, singularities corresponding to *complex* values of the expansion parameter may limit the convergence of the series, but these singularities may have no direct physical "meaning".

A three-step hybrid analysis technique, which successively uses the regular perturbation expansion method, the Pade expansion method, and then a type of Galerkin approximation, is being developed and studied. Currently, it is being applied to several model problems which develop boundary layers as a certain parameter becomes large. These problems involve ODE's, PDE's, and integral equations. In particular, the technique appears to simulate these boundary layers by producing approximate solutions with real or complex singularities which lie just outside the domain of interest.

The hybrid Pade Galerkin technique will be applied to and studied in the context of some ODE's and elliptic PDE's which develop internal layers, boundary layers, and two classes of integral equations of the first kind.

3

DAVID GOTTLIEB

Numerical solutions of PDE's

Our long term research objective is to develop spectral algorithms general enough to handle complex geometries. Both single domain and multidomain approaches are considered.

Unstructured meshes have long been recognized as a means to grid arbitrary geometries. We have therefore begun to develop spectral methods for unstructured grids. In an ICASE report with M. Carpenter we have outlined the application of spectral methods on one dimensional unstructured grids. We are now extending this idea to complex shape domains in two dimensions. This is a one domain approach rather than multidomain method.

In addition, we have derived a method to determine the optimal number of subdomains in domain decomposition techniques for hyperbolic equation.

Suppose we wish to resolve certain wavelike phenomena in two space dimensions. Let us assume that we have k such waves (the complexity of the problem increases with k), and that we need to resolve these to within a given tolerance, $e^{-\epsilon}$ say, on a parallel machine with P processors. The multidomain discretization involves decomposing the 2D domain into Q subdomains of N points each. Our formula dictates that $N = (\epsilon z)^2$, $Q = \left(\frac{\pi k}{\epsilon z} e^{\frac{1}{z}}\right)^2$ where z satisfies the equation $z(z-2)e^{\frac{1}{z}} = \sqrt{P} \frac{2t_b}{At_a} \frac{1}{\pi k \epsilon}$, where t_a is the time for an operation within the processors and t_b is the interprocessor communication.

MAX GUNZBURGER

Optimal control of fluid flows with application to transition delay

The systematic suppression of instabilities in incompressible boundary layer flows is considered as an optimal control problem. The goal is to determine optimal strategies for suppressing these instabilities without having to inject ad-hoc or empirical information and without interference from the designer.

Optimal control methodologies are developed and applied to this problem. The specific control mechanism considered is the injection or suction of fluid through orifices located along a small portion of the bounding wall. The objective is to minimize the shear stress along another small portion (downstream from the control orifices) of the bounding wall. The methodology involves the application of the Lagrange multiplier technique to derive a system of partial differential equations from which optimal controls may be determined. This system consists of the Navier-Stokes equations coupled with an adjoint system of equations for the co-state or Lagrange multiplier variables and an optimality condition that relates the optimal control to the state and co-state variables. Appropriate boundary, initial, and terminal conditions are also defined. The results demonstrate that instability suppression can be achieved without any a priori knowledge of the disturbance, e.g., of its frequencies, phases, type, or amplitudes.

Future plans include the use of other control mechanisms such as shape controls, different objective functionals, and the application of the optimal control methodologies to more complex flow configurations. This research was conducted in collaboration with Ron Joslin, R. Nicolaides, Gordon Erlebacher, and M. Y. Hussaini.

MAURICE HOLT

Review of Godunov methods

Godunov's method for solving unsteady problems in Gas Dynamics is described and discussed at length in Holt "Numerical Methods in Fluid Dynamics" (1984). In the succeeding decade several extensions of the method have been proposed which increase its accuracy to second order while retaining the properties in the original method of *monoticity* and absence of oscillations in *shock capturing*. The contributions most clearly related to the first Godunov scheme, notably by Colella and Woodward, Roe and van Leer are reviewed and added to the earlier account. At the time of writing Holt (1984), these contributions either had not been completed or not yet recognized as Godunov extensions. The revised chapter in my book (to appear as a 3rd edition) will include these extensions and the present report is a preliminary version of this coverage.

The extensions treated specifically are the MUSCL (Monotonic Upstreamcentered Scheme for Conservation Laws) scheme of van Leer (and Woodward), the Piecewise Parabolic Method (PPM) of Colella and Woodward and the characteristics based scheme of Roe. The report first deals with these methods as originally presented. Thus, PPM is described only in application to the unidirectional wave equation, with indication of its extension to the Eulerian gas dynamic equations. The MUSCL scheme as developed by van Leer, is applied to the one-dimensional Lagrangian equations. Roe's scheme, freer than the other schemes of algebraic details, is presented in general form.

In the final version of this review the TVD and ENO schemes initiated by Harten will be covered. Also the constraint needed to maintain the monotonic property of higher order Godunov schemes has been developed in better and more convenient form by Huynh. This is based on the simplification of the second order upwind scheme in MUSCL by Hancock and these improvements will be included.

ANGELO IOLLO

Optimal shape design for Euler equations

A classical problem in engineering is to define the shape of a vehicle to achieve a required performance. In fluid dynamics, techniques have been developed to solve the following inverse problem: given a pressure or a velocity distribution over an aerodynamic body, determine the corresponding geometry. A broader category of problems can be solved by means of optimization, provided that one is ready to accept the necessity of computing the flow field hundreds of times. In using models of increased complexity to describe the flow field, the development of new algorithms is necessary in many cases to reduce the computational load. We investigated one method for achieving this reduction. A functional or cost function is determined such that its minimum represents an optimal solution. By introducing a set of Lagrange multipliers, the gradient of the functional can be calculated with respect to the geometry by computing the flow field only once for each gradient evaluation. We studied a method for calculating the conditions that the Lagrange multipliers must satisfy at the boundaries and at the shock. Results were obtained for two-dimensional internal flows with embedded shocks.

Still this method is not practical in some cases, since it is necessary to compute the flow field and the costate equations many times to reach the minimum. Furthermore, in general the number of iterations to achieve the minimum grows more than linearly with the number of controls used.

Ta'asan proposed an algorithm to reduce the cost of the optimization to that of a single analysis. The idea is to solve the flow equations, the costate equations and the optimality condition at the same time. Further work includes research to apply this idea to Euler equations.

This research was conducted jointly with M.D. Salas of NASA Langley and S. Ta'asan of Carnegie Mellon University.

LELAND JAMESON

Applications of WOFD

The objective is to extend the Wavelet Optimized Finite Difference (WOFD) to higher dimensions and to a variety of applications. The motivation is to be able to solve problems which can not be solved by conventional numerical methods due to computer memory and CPU time limitations. WOFD excels when the application is such that the scale size of physics is different in different regions of the domain. That is, if there is a shock or some other type of small scale phenomenon in a small portion of the domain while the remainder of the domain is dominated by large scale physics, then WOFD can usually greatly reduce the amount of memory and CPU-time required.

To date, all applications to one-dimensional problems have proved successful. This includes applications to the ICASE equation with M. Carpenter in which we tested WOFD and Carpenter's provably stable and very promising adaptive grid numerical method. In addition, WOFD has been applied to the inviscid Burger's equation with R.B. Bauer and again the results were very good. In a project with T. Jackson, WOFD was applied to problems in 1 dimension combustion, again with success.

The future will consist of extending WOFD to higher dimensions and more interesting problems. The program is completely adaptive in all dimensions and is, therefore, difficult to write and debug. Hence, a working two-dimensional version of the program has been written and is operational. Eventually, a three-dimensional version will be written.

DAVID A. KOPRIVA

Spectral multidomain methods for compressible flows

Most spectral multidomain methods are complicated to implement in multiple dimensions, due to the special treatment of corner and edge nodes. Our objective is to generalize current multidomain algorithms to allow arbitrary polynomial orders in each subdomain while simplifying the implementation. We have continued the development of a new family of multidomain spectral methods for the solution of compressible flow problems that is based on a fully staggered grid. The solutions are defined at the nodes of a Gauss quadrature rule. The fluxes are evaluated at the nodes of a Gauss-Lobatto rule. The fluxes are needed at interfaces, and an approximate Riemann solver is used to determine the characteristic decomposition. In two space dimensions, the method does not include the corners of subdomains, making the method easy to use to compute complex grid

6

topologies. The first version of the method required that the approximation be conforming. This meant that subdomains were required to intersect completely along a side, and that the polynomial approximations along the side were of the same order. In other words, the grid lines from each subdomain must meet at an interface. This requirement limits the flexibility to do local refinement of the subdomains and polynomial order.

We have now developed a non-conforming staggered-grid approximation that allows local refinement of subdomains and polynomial order. Subdomains can be divided in two, and it is not necessary to have the same polynomial orders on either side of an interface. To accommodate the different approximations, a least-squares projection is made onto a "mortar" along which the Riemann problem is solved. The result is then projected back onto the contributing subdomains. This mortar procedure retains the global conservation of the original conforming staggered-grid approximation. Numerical studies show that the approximation is exponentially convergent.

The next stage is to extend the staggered grid method to the advection-diffusion equation and the compressible Navier-Stokes equations. A method for treating the advective fluxes has been devised. It has been successfully implemented on one-dimensional linear, Burgers' and Navier-Stokes problems. The development of two-dimensional applications is now underway.

ROBERTO MARSILIO

A numerical method for solving PNS

In the study of three-dimensional supersonic flows it is common to be faced with complex interactions involving shock waves and viscous layers. Such interactions often produce dramatic changes in the flowfield features. Shock-induced separations of the viscous layer make the flow vortex-dominated close to the wall. Vortical structures are also likely to appear, for sufficiently high Reynolds numbers, in zones where convective effects are predominant, due to the instability of the slip surfaces resulting from shock/shock interactions.

Numerical tests are usually helpful in the investigations of such complicated fluid patterns, as they provide a tool to observe and possibly to understand the origin and the effects of the numerous phenomena triggered by shock/shock and shock/viscous layer interactions. It is clear, however, that in order to obtain good numerical results comparable with experimental data, a good description of the flowfield is necessary, especially when multiple vortical structures are present.

The only completely correct way of solving numerical three-dimensional compressible viscous flows is to integrate in time the full Navier-Stokes equations until a steady-state, if it exists, is reached. Although this approach is sometimes affordable, if many grid points are needed for detailed resolution of complex flow features, it could be excessively time and memory consuming. In the case of supersonic steady-state flows, however, this practical difficulty can be partially circumvented with the aid of the approximate form of the full Navier-Stokes Equations known as *Parabolized Naviers-Stokes Equations*, (PNS).

In the present approach, the governing equations are integrated explicitly and the physical domain is discretized according to a finite-volume technique. The convective part of the equations (inviscid fluxes) is treated following a flux-difference-splitting method with an approximate solution of a Riemann problem at each cell interface, while diffusive terms (viscous fluxes) are calculated using a central scheme. Second order accuracy is reached by means of an ENO scheme with linear

reconstruction of the solution at each integration step. Presently, only inert gases in the laminar regime have been considered, but a future extension to thermochemical or turbulence models is certainly possible.

Interaction of Homogeneous Turbulence with a 3D Shock Wave

A fundamental understanding of compressible turbulence is necessary for the development of supersonic transport aircraft. Compressibility effects on turbulence were found to be significant when the energy associated with the dilation fluctuations is large or when the mean flow is significantly distorted (expanded or compressed). The presence of shock waves is an important feature that distinguishes high-speed from low-speed flows. Understanding the interaction of homogeneous turbulence with a shock wave has both general interest and fundamental importance in understanding the interactions of turbulent boundary layers with shock waves which occur in many practical engineering applications: the flow inside a high speed compressor or a gas turbine, the flow over wings in supersonic aircraft, and the intake flow to a supersonic ramjet.

Numerical simulation using turbulence models is becoming a standard tool in aerospace technology. Most current models of compressible turbulence are based on incompressible concepts. A better understanding of the underlying physics could lead to improvements in turbulence models, leading to more efficient designs. There is, therefore, a need to understand compressible turbulence.

The numerical approach is based on the time-dependent integration of the fully three-dimensional Navier-Stokes equations. The governing equations written in the quasi-linear form where the spatial derivatives are computed using a 6^{th} order compact scheme while the time integration is performed by a 4^{th} order Runge-Kutta scheme. The shock is fitted using a three-dimensional shock-fitting procedure in which the shock is treated as moving internal boundary separating two (upstream and downstream) regions. To simulate shock turbulence interaction, we initialize the upstream flow (supersonic part) by superimposing homogeneous turbulence on the mean flow.

In future work we will validate the code by comparing the results obtained with appropriate analytical/numerical solutions; part of the work has just been completed for the two-dimensional configurations. The final goal of this work is to better understand the behavior of the turbulence with particular regard to the shock/turbulence interaction. This research was conducted in collaboration with Gordon Erlebacher and M. Y. Hussaini.

DIMITRI MAVRIPLIS

Unstructured multigrid techniques for the Navier-Stokes equations

The overall objective of this work is to develop efficient solution procedures for the Navier-Stokes equations on unstructured meshes with low memory overhead and with eventual application to parallel machines.

Multigrid methods offer a promising approach to achieving rapid convergence of the Navier-Stokes equations without incurring large memory overhead. For unstructured meshes, the construction of coarse mesh levels is non-trivial. A first approach operates on a set of non-nested coarse and fine meshes, and transfers variables between the various meshes of the sequence using linear interpolation. This non-nested multigrid approach has been demonstrated for the three-dimensional Reynolds-averaged Navier-Stokes equations using a single field equation turbulence

8

model on high-lift swept wing configurations with very fine grids of 2 million points or 10 million tetrahedra. Convergence to steady state for such cases can be obtained in 200 to 400 multigrid cycles, which can be performed in 1.25 hours using all 16 processors of the CRAY-C90.

An alternate approach uses agglomeration to construct coarse grid levels. This approach has the advantage that the coarse levels are constructed automatically using a graph-based algorithm. However, the coarse levels are no longer simplicial meshes and alternate discretization procedures must be employed to construct the coarse grid problems. The use of Galerkin coarse grid operators for Navier-Stokes problems has been investigated. The agglomeration multigrid algorithm using this construction has been implemented for three-dimensional Navier-Stokes problems. The solution of viscous turbulent flow over a partial span wing-flap geometry has been performed using a fine grid of 2.3 million points and 13.6 million cells. The code was parallelized on the CRAY-C90 and achieved a 5 order reduction in the residuals over 300 multigrid cycles, which required 1.25 hours of wall clock time using all 16 processors of the CRAY-C90. These results were presented in a paper co-authored with V. Venkatakrishnan at the 33rd AIAA Aerospace conference in Reno Nevada, in January 1995.

In related work with Eric Morano and V. Venkatakrishnan, stencil based coarsening strategies have been investigated for constructing coarse overset meshes for traditional unstructured multigrid methods for anisotropic problems. The results indicate that such techniques can substantially enhance multigrid convergence rates for Laplace's equation discretized on anisotropic meshes.

Future work will center on extending the stencil based coarsening strategies to agglomeration multigrid methods and to the solution of the Navier-Stokes equations, first in two dimensions and subsequently in three dimensions.

ERIC MORANO

Coarsening strategies for unstructured multigrid techniques with application to viscous flows

When solving the equations that represent viscous flows, such as the Navier-Stokes equations, it is important to take into account the presence of a boundary layer and/or other phenomena associated with the flow such as vortices. Correct discretization of the equations on the domain to be studied requires a very fine and highly stretched grid. The use of such grids leads to extremely long computation times. While multigrid methods can be used to accelerate the convergence for such problems, they have limitations when used with highly stretched grids; the gain in performance is generally much lower than that achieved for inviscid flows and difficulties remain in the generation of the different coarse grids. The purpose of this work is to analyze and to overcome the loss of performance of the multigrid scheme and to propose a simpler and faster way to produce grids of good quality in order to restore an "ideal" convergence rate.

This research was conducted in collaboration with D. Mavriplis and V. Venkatakrishnan.

The Laplace's equation

The multigrid (MG) technique has been demonstrated as an efficient technique for solving inviscid flow problems . However, for viscous flow problems, convergence generally degrades. This is due to the use of stretched meshes with aspect-ratio AR $= \Delta y / \Delta x \ll 1$ required in order to capture the boundary layer near the body. The problem addressed here concerns the improvement

of the convergence of the MG algorithm through the generation of the coarse grids. The work focuses on the solution of Laplace's equation on stretched grids, since the convergence degradation of the Navier-Stokes equations on such meshes is reproduced by this simpler equation. The discretization of the equation, through a Galerkin finite-element formulation, can be written as a sum over edges:

$$\int_\Omega \Delta u \; \varphi \; dxdy = \sum_{edges} c_{edge}(u_j - u_i)$$

where Ω represents the domain and φ is a standard test function of compact support. On a stretched mesh, the distribution of nodes in the stretching direction represents the low frequencies of the signal, whereas the distribution of nodes in the direction normal to the stretching represents the high frequencies. Because of the nature of the smoothers commonly used, the MG scheme efficiently damps the high frequencies. Hence the idea of semi-coarsening in the direction normal to the stretching.

An MG algorithm that performs properly must provide linear-type mesh-independent convergence. In order to do so, the MG scheme must satisfy two essential properties. The smoothing property, which relies on the smoother and the ellipticity of the equation, is here assumed to be true. The approximation property provides the relative interpolation error of the solution on two different grids, which is satisfied when the grids are nested. This is not always the case when the MG algorithm is applied to unstructured grids.

The usual full-coarsening method consists in removing every second vertex in each direction on a structured grid, which results in a decrease in the number of nodes of the coarse grid by a factor 4. The complexity of a V-cycle tends to 4/3 WU (i.e. a Work Unit corresponds to the computation of one residual on the fine grid). By contrast, the semi-coarsening technique produces an overall complexity that tends to 2. The value of the coefficient associated with each edge in the formula above is determined by the geometry of the surrounding elements (triangles). The smaller the length of the edge the larger the value of the coefficient. The present semi-coarsening algorithm selects a node that remains on the coarse grid and removes its connecting neighbor that corresponds to the edge associated with the stiffest coefficient. The reconnection of the remaining set of nodes is performed through a variant of the Delaunay triangulation technique. In order to validate this concept, a sequence of test-cases is performed. The selected smoother is the weighted Jacobi relaxation scheme. The MG convergence histories, depicting the logarithm of the normalized residual on the finest grid versus the number of cycles, are analyzed. The transfer operators are linear. Dirichlet boundary conditions define the solution at the extremities of the mesh.

The non-stretched (AR = 1) fully-nested test-case serves as the standard. A solution is obtained after 12 cycles which corresponds to a convergence rate of 0.15. In general, the full-coarsening technique applied to stretched meshes results in convergence rates which are approximately equal to 0.80. Since this is considerably different from the standard linear convergence, it cannot be considered mesh-independent. The semi-coarsening technique, on the other hand, results in a mesh-independent convergence with a rate equal to 0.15. A fully unstructured MG experiment has also been performed on a fine grid comprising 19366 vertices. The average minimum aspect-ratio is equal to 3.7×10^{-5}. A sequence of 7 grids, which results in a W-cycle complexity equal to 12.5 WU, provides a linear type of convergence with a rate equal to 0.23.

In summary, in this work a new coarsening algorithm relying on the discretization of the equation has been introduced which enables flexible applications. This results in convergence rates for highly-stretched unstructured meshes similar to those for standard cartesian structured non-stretched meshes. The MG convergence demonstrated linear, hence mesh-independent rates.

JAMES J. QUIRK

Numerical investigation of pulsating detonations

From experimental observations it has long been known that detonation waves can exhibit a one-dimensional pulsating instability, the so-called galloping instability, which can produce pressure fluctuations which are several times larger than the corresponding stable detonation pressure. We have conducted a detailed numerical investigation of this phenomenon for a one-dimensional detonation wave driven by a chain-branching reaction model representative of a Hydrogen-Oxygen system. This study serves two main purposes. Firstly, it sheds light on a phenomenon that is not well understood. Secondly, it provides a rigorous validation test for any numerical method developed to investigate multi-dimensional detonation phenomena such as will be needed for the design of an oblique detonation wave engine.

Although our study is only one-dimensional, a galloping instability involves a tremendous range of physical scales which must all be resolved in order to produce reliable results and a one-dimensional calculation can therefore be surprisingly expensive to compute. Failure to resolve the smallest scales can lead to completely erroneous physical behavior and will not even give qualitatively correct predictions. Therefore we made use of a sophisticated mesh refinement algorithm to ensure that all the pertinent physical scales were properly resolved and all results were checked to see that they were grid converged. Additionally, we were able to use the results from a linear stability analysis to check the accuracy of certain predictions. Thus for those cases of interest where the analysis is known to be invalid, we have confidence that our numerical predictions are reliable. This investigation, which has now been written up and will appear as an ICASE report, has revealed several hitherto unknown features of galloping instabilities.

The present study forms a self-contained piece of work. However it does provide a springboard from which to investigate certain related phenomena such as detonation quenching and several options for further work are currently being considered. This research was conducted in collaboration with Dr. Mark Short.

PHILIP ROE

Multidimensional algorithm development

Decomposition of the Euler equations and other multidimensional systems of PDEs into canonical form appears to be the most faithful analogue of the characteristic decomposition that is so successful when the equations involve only one space dimension and time. In joint work with Eli Turkel and my student Lisa Mesaros (ICASE Student Visitor, 1994) we aim to explore the forms that are available, and their implications for numerical computation.

For example, in two dimensions, the regular case is that an $N \times N$ system can be split into $N - 2n$ ODEs and n Cauchy-Riemann systems, but there are cases where this splitting is singular, for example, at stagnation points and sonic lines for the Euler equations. Current work concentrates on the implications of these singularities.

Bicharacteristic schemes for linear wave propagation

Following on the work described in ICASE Report 94-65, which developed a scheme for long-range wave propagation, free from numerical dissipation but correctly oriented to carry characteristic information, we are attempting to increase the accuracy of the method from second order to fourth order. The extension is easy in one space dimension and results in schemes that are compact, thus preserving full accuracy up to the boundary, and that maintain phase and group velocities to within 1% with only two or three points per wavelength.

In higher dimensions, careful selection of the stencil and data structure seem to be required in order to achieve stability, and only partial results are currently available. If successful, however, the schemes could have substantial impact on the calculation of electromagnetic and anisotropic elastic waves.

CHI-WANG SHU

Shock capturing numerical methods

Methods which are of high order in the sense that local truncation errors are small when the flow is smooth can usually give increased resolution at lower cost than lower order methods, even for problems with shocks. We study and apply such methods in the class of finite difference, finite elements and spectral schemes.

Jointly with Harold Atkins at NASA Langley, we are continuing to analyze and apply the discontinuous Galerkin finite element method. Currently we are investigating efficient ways to treat the nonlinear terms in the integration. Jointly with Gordon Erlebacher and Yousuff Hussaini, we are continuing our investigation of shock vortex interaction problems. Results of shock interaction with strong longitudinal vortex, where shock bifurcation and secondary shocks appear, are simulated by ENO method and analyzed physically.

Research will be continued for high order methods in finite difference, finite elements and spectral schemes.

DAVID SIDILKOVER

Unification of some advection schemes in two dimensions

Two classes of the genuinely multidimensional nonlinear positive advection schemes are known. One of them is of the control-volume type. Its novel feature is the smoothness monitor function (limiter) that compares gradients in *different* directions. This leads to a high-resolution advection scheme with a minimally-enlarged stencil, and with good multigrid capabilities. Another class was the "fluctuation-splitting" schemes. These schemes are intended to be coded as a loop over triangular elements, employing no data external to the triangle. Several non-linear variants of this technique have been devised that preserve positivity.

The main objective of this work is to formulate a "unified approach" which includes the desirable properties of both classes. A strong relationship between the two classes of the advection schemes was established. This allowed formulation of a "unified approach" - a class of schemes having the following properties:

- the simplicity of dealing with unstructured meshes as in case of the fluctuation-splitting schemes

- the use of two-dimensional limiter, which allows artificial compression and achieves good accuracy (linearity preserving property) for both the inhomogeneous and homogeneous cases while preserving the positivity property.

This work was done in collaboration with P. L. Roe from the University of Michigan.

Future plans include the design of new limiter functions which will allow further accuracy improvements.

Genuinely multidimensional schemes for the compressible Euler equations

The key advantage of the recently constructed genuinely two-dimensional schemes for the compressible Euler equations is that the Gauss-Seidel relaxation, a simple and very efficient smoother can be applied directly to the resulting discrete equations. This allows construction of simple simple, robust and efficient steady-state multigrid solvers.

Further verification of the performance of the algorithm based on this scheme were conducted for various flow regimes. The two-dimensional scheme was formulated in a form suitable for unstructured grids. This genuinely multidimensional approach was generalized for the three-dimensional case, both for the scalar advection equation and the Euler systems of gas dynamics.

Future plans include an unstructured grid implementation and verifying the performance of the constructed scheme when applied to the real-world problems. This work will be done in collaboration with W. Kyle Anderson of NASA Langley.

RALPH SMITH

Implementation of PDE-based controllers

Throughout the last several years, PDE-based control methods for structural and structural acoustic systems have been developed. While numerical tests have demonstrated the efficacy of the methods, few experiments concerning the implementation of these methods have been performed. In collaboration with D.E. Brown, V.L. Metcalf, R.J. Silcox (Acoustics Division, LaRC) and Y. Wang (Brooks Air Force Base), issues concerning the implementation of these methods are being investigated.

Recent research has centered around the use of PDE-based control methods to reduce vibrations in a circular plate. An advantage of PDE-based methods over standard frequency response input/output techniques is the capability to directly control both transient and steady-state responses. In initial experiments, a centered hammer impact was used to excite the plate and integrated data from an off-center accelerometer was used to reconstruct the state. LQG control theory was used to calculate controlling voltages to a centered piezoceramic patch bonded to the plate. In

a second set of experiments, a periodic steady-state voltage was applied to a secondary patch and results demonstrating the control capabilities in the presence of both transient and steady-state vibrations were obtained. In both cases, attentuation levels on the order of $80 - 90\%$ demonstrate the effectiveness of the PDE-based controller for this plate system and indicate the potential for applying the methods in other structural and structural acoustic systems.

Upon completion of the experiments involving the circular plate, we will begin experiments with a hardwalled structural acoustic system with a vibrating plate at one end. We will also continue extending the analysis and numerical techniques to models in which the structure enclosing the acoustic cavity is a vibrating shell. This research was conducted in collaboration with H. T. Banks.

The modeling of inexact boundary conditions

A problem common to many structural and structural acoustic applications is the modeling of inexact boundary conditions. Even when heavy boundary clamps are employed, energy is often lost at the boundary, leading to natural frequencies lower than those predicted by models in which truly clamped boundary conditions (zero displacement and slope) are assumed. The modeling of "almost clamped" boundary conditions to account for this boundary energy loss is being investigated in collaboration with Y. Wang (Brooks Air Force Base).

A model for the "almost clamped" boundary conditions has been developed in which slight displacement and slope variations are permitted with the degree of variation governed by parameters to be estimated through fit-to-data techniques. The well-posedness of the model has been demonstrated and numerical studies concerning the estimation of parameters in the model and the usual clamped-end model limit have been performed.

We are currently setting up experiments with a circular plate having loosened boundary clamps. With data from these experiments, we will use parameter estimation techniques to evaluate the ability of the model to account for the physics of the system. This will provide a step toward the modeling of inexact boundary conditions in more complex structures. This research was conducted in collaboration with H. T. Banks.

SHLOMO TA'ASAN

Pseudo-time methods for constrained optimization problems

The pseudo-time one-shot methods for optimization problem is an approach aimed at solving PDE constrained optimization problems with a cost comparable to that of solving the PDE just a few times. This direction emerges as a result of examining the drawbacks of existing methods including the one shot multigrid method. The latter, although very efficient, requires multigrid structure and is not easily implemented.

The main idea is to regard the solution methods for constrained optimization problems as discretization of some pseudo-time evolution problem. Methods based on gradient information can also be viewed in this way; however, they are easily shown to be excessively complex. The minimum for the optimization problem is viewed as an intersection point of hypersurfaces described by the state equation, the costate equations and the design equation. Gradient based methods are marching techniques on the intersection of the state and costate hypersurfaces, and require the solution of two PDE per step. In contrast, the new methods are marching on the design

hypersurfaces and require a solution of a much smaller problem per each step, namely, another boundary condition. The convergence of the new methods is checked using energy estimates for the full system. Results for several model problems involving PDE in two dimensions have been obtained. The efficiency obtained in all examples shows that at a cost of 2-3 times the cost of solution of the original PDE, one reaches the solution of the full optimization problem.

Future theoretical work involves a more sound basis for the new methods, including understanding the relation between the character of the minimum and the convergence rate. Applications will be carried out in design of airfoils with Euler and Navier-Stokes equations.

Pseudo-time methods for aerodynamics design problems

Optimization problems governed by the fluid dynamics equations are very challenging and require very efficient algorithms; the solution time for even a single analysis is of the order of a few CRAY hours for the Navier-Stokes equations.

Application of pseudo-time methods for aerodynamics optimization application were carried out together with G. Kuruvila and M.D. Salas and A. Iollo. Experiment with the small disturbance equation in airfoil design showed that full design is done at a cost of 3-4 analyses. The method used a Gauss-Seidel relaxation for converging the state and costate equation, where at each time step the boundary conditions as well as the design equations have been satisfied. Similar behavior was obtained for the Full Potential equation and the design included lifting airfoils. Further application of the ideas have been carried out using the Euler equations in designing nozzles in two dimension.

The extension of ideas to handle lifting airfoil design governed by the Euler equations will be a subject of future research. Further application will include the Navier-Stokes in two and three dimensions. The inclusion of algebraic constrains together with the PDE will also be studied.

V. VENKATAKRISHNAN

Solution techniques for unsteady flow simulations over complex geometries

Work is continuing in devising efficient and accurate methods for computing unsteady flows over complex geometries. The implicit method allows arbitrarily large time steps. Following a backward difference time discretization, the nonlinear system at each time step is solved by using the agglomeration multigrid procedure. This work is being done jointly with Dimitri Mavriplis.

It has been verified by numerical experiments that the "mass matrix" may be lumped for schemes possessing second order spatial accuracy or less without suffering any adverse consequences for unsteady problems. If a control volume viewpoint is adopted, the mass matrix is an important issue on nonuniform grids, where the cell vertex and the centroid of the control volume do not coincide. Local analysis would reveal a drop in the order of accuracy to first order. The numerical experiments on random grids indicate that this does not happen. Second order accuracy is maintained even with the mass lumping. However, for higher order accurate schemes, the mass matrix appears to be crucial. For a third order accurate spatial discretization achieved by using a quadratic reconstruction procedure, mass lumping degrades the accuracy to second order, whereas employing a finite volume mass matrix restores the third order accuracy. The finite volume mass matrix is defined as the average of the reconstruction polynomial and differs from the definition employed in the finite element method. In these experiments, time integration is performed using

15

a Runge-Kutta method of equal or of higher order accuracy than that of the spatial discretization. Two-dimensional unsteady Euler computations about oscillating airfoils have also verified that the accuracy of the scheme is not compromised by mass lumping.

Unsteady laminar viscous flows have also been solved by the above methodology. The start-up flow over a cylinder has been computed and results have been compared with experiments. In order to compute flows over bodies in relative motion, a procedure has been devised so that the validity of the grid is preserved. It allows for large scale relative motion between bodies, such as the flaps and the main airfoil. The procedure involves the use of smoothing, a spring analogy for moving grid points and swapping of edges to preserve the Delaunay triangulation. The methodology is currently being integrated into the unsteady flow solver.

A survey paper on unstructured grid flow solvers was presented as an invited paper at the AIAA 33rd Annual Aerospace Sciences meeting held in Reno, NV in January. The paper reviewed the developments in the following areas: finite volume and finite element spatial discretizations, explicit and implicit schemes, adaptation and parallel computing.

A set of three lectures was delivered at the von Karman Institute for Fluid Dynamics, Brussels, Belgium in March. It was accompanied by prepared notes. The lectures dealt with the development of implicit schemes on unstructured grids for steady flows, explicit and implicit schemes for unsteady flows and parallel computing aspects, including grid partitioning, flow solvers and grid adaptation.

FLUID MECHANICS

PONNAMPALAM BALAKUMAR

Non-linear equilibrium solutions in Couette flow

Plane Couette flow serves as a simplified model of a boundary layer, and is simpler to analyze both theoretically and computationality. Thus, insight into the dynamics of boundary layers can be gained by enhancing our knowledge of the model flow. Unfortunately, plane Couette flow is stable at all Reynolds numbers to linear disturbances. Therefore, we investigate the non-linear stability of plane Couette flow.

Plane Poiseuille flow is unstable to small disturbances for Reynolds numbers exceeding 5771. We find that this flow is subcritically unstable at finite amplitudes. Calculations show that the non-linear equilibrium solutions exist for Reynolds numbers up to 2950. Starting with these non-linear equilibrium solutions for the Poiseuille flow, we gradually increase the velocity of the upper wall to find the equilibrium solutions for plane Couette flow. However, when the velocity of the wall increases, the convergence to the equilibrium solution deteriorates and one needs to include a large number of modes.

In the future, we plan to investigate the non-linear equilibrium solutions in two-dimensional flows.

FRANCOISE BATAILLE

Numerical simulation of compressible turbulence

We investigate compressibility effects on homeogeneous and isotropic turbulence. Using two-point closures (E.D.Q.N.M. for Eddy Damped Quasi Normal Markovian) in Fourier space, we find that the compressible spectrum scales as $K^{-5/3}$ in the inertial zone for short times and has an asymptotic state of $K^{-11/3}$. The aim of the study is to find out whether this behavior can be found using direct numerical simulation (DNS).

The simulations are done using the simulation code developed by G. Erlebacher. Since results are difficult to obtain with DNS, the study is done using Large Eddy Simulation. Different initial conditions were tested by changing the shape of the initial incompressible spectrum. Both analytical spectra and the Comte-Bellot/Corrsin spectrum were used. Initial results show that the $K^{-5/3}$ and $K^{-11/3}$ behavior can be recovered. This is especially true for particular initial conditions which already embody the correct spectrum slope information. Both the Smagorinsky and the dynamic model have been implemented. The same tests were performed with these two models and similar results were obtained. Although more tests are necessary, it does not appear that a more complex subgrid-scale model is required for this type of turbulence (isotropic).

The next step of this study will be to confirm the results of the compressible spectrum behavior and to study their sensitivity to the turbulent Mach number.

ALVIN BAYLISS

Response of an array of aircraft panels to jet noise

We study numerically the response of an array of aircraft type panels to jet noise. In our model the panel response and radiation are fully coupled to the Euler equations governing the generation and propagation of the jet noise. We consider a model which includes the jet domain, where the noise is generated, panel assembly and a radiation domain (the analogue of the aircraft interior). The primary objective is to study panel response to jet noise so as to be able to predict and control panel response and so minimize structural fatigue and interior noise levels. The primary focus of current research is to incorporate forward motion effects into the model (i.e., to model a jet in flight). The objective is to determine the role played by forward motion on installation effects from the nearby panel array and on the response and the acoustic radiation from the array.

A detailed analysis of the effect of uniform forward motion for a subsonic jet has been accomplished. The results include detailed computation of the radiation pattern from the panels for a jet in forward motion. An ICASE report has been prepared and this work will be presented at the CEAS/AIAA 1st Joint Aeroacoustics Conference in June, 1995. A second paper on this work is being prepared and will be submitted shortly.

Future work will include the incorporation of boundary layer effects within the forward motion model, so that the effect of wall boundaries on panel response and radiation can be simulated. Further work will include the study of accelerated flow and the extensions of this work to supersonic jets. This research is done in collaboration with L. Maestrello, C. C. Fenno, Jr., and J. L. McGreevy.

AYODEJI O. DEMUREN

Near-wall model based on RNG theory

Turbulence models are usually derived on the assumption of high Reynolds numbers, which permits significant simplification and the neglect of direct viscous effects. However, in the immediate vicinity of walls, viscous effects are important and Reynolds numbers are low. The turbulence models must be modified in this region, but most existing modifications are ad-hoc and contain little physical reasoning. They are merely correlations of experimental data. Not surprisingly, they fail in computations of flows which differ significantly from those used to calibrate them. Renormalization Group Theory (RNG) presents a sounder basis for developing near-wall models.

In collaboration with Ye Zhou, a near-wall k-e model is proposed based on recursive RNG theory. The original high Reynolds number version of the model was derived based on the assumption of the Kolmogorov $k^{-5/3}$ law. At low Reynolds numbers this law may not exist and the energy spectrum may be dominated by the dissipation range. Based on a general law, which accounts correctly for the dissipation range, modifications are derived for the $k - \epsilon$ model which automatically accounts for low Reynolds number effects.

The model is now being applied to a series of simple near-wall flows such as boundary layers in zero-, adverse- and favorable- pressure- gradients and plane channel flows.

GORDON ERLEBACHER

Numerical simulations of shock/vortex interaction

Fundamental understanding of shock-turbulence interaction, both from the point of view of acoustics, and that of turbulence modeling remains a primary research target. It is known from large scale Direct Numerical Simulations (DNS) of the compressible Navier-Stokes equations that shocks amplify the vortical component of turbulence, and generated acoustic pressure waves. Tests against linear theory have partially explained the observed phenomena. We have applied our 2-D shock-fitted compressible Euler code to the study of shock/turbulence and shock/vortex interaction with a view towards quantifying the processes which come into play when the flow parameters induce nonlinear effects. The information gained from these calculations will help the three-dimensional simulations that will follow, as well as the researchers developing new turbulence models.

Using a 6^{th} order compact scheme in space, and a 4^{th} order time advancement scheme, we have applied our shock-fitted code in two ways. First, we have generated random solenoidal velocity fields (of varying intensity) upstream of the shock, and advected it through the shock. Comparisons with linear theory have determined that a 1% turbulence intensity (with respect to a $M = 2$ upstream mean flow) is sufficient to engender nonlinear effects downstream of the shock (the turbulent evolves from an axisymmetric state at the shock to an isotropic state far downstream). Further checks were performed using a linearized version of the nonlinear code. A second problem involves the passage of several isentropic vortices (of varying circulation) through shocks of varying strength. Budgets of vorticity, dilatation and kinetic energy allow us to quantify the nonlinear sources of energy production. The analysis is in progress.

To investigate detonation-turbulence and detonation-vortex interactions, the code will be modified to included reaction via a simple one-step heat release mechanism with a finite-length reaction zone.

SHARATH S. GIRIMAJI

Direct numerical simulations of scalar mixing and utilization of data to validate LES scalar dissipation model

Direct numerical simulations of passive scalar mixing in forced isotropic turbulence were performed to study the spectral behavior of the scalar field during mixing. The resulting DNS data was used to examine many of the assumptions routinely made during model derivation. The findings suggest that some of the previously held notions about the spectral behavior of a scalar field need to be modified. This work has culminated in a reasonably well validated model expression of the large eddy scalar dissipation. However, further testing and improvements are still necessary.

Investigation of difference between the original and modified restricted Euler equation

Calculations were performed to closely examine the differences and similarities between the original and restricted Euler equations. The results clearly indicate that the modified restricted Euler equations are more accurate than the original version for general homogeneous flows. It was also demonstrated that the modified restricted Euler equation are capable of capturing extrinsic orientation properties of fluctuating velocity gradients. Proposed future work in this area includes

the investigation of the effect coordinate frame rotation on turbulence, an issue of great practical importance.

Investigation of inverse energy cascade in Burgers turbulence

Direct numerical simulations of Burgers equation are performed to examine if Burgers turbulence permits an inverse cascade of energy from small to large scales. At this stage, we have shown that inverse energy transfer is possible. However, it is not clear if this transfer is in the form of a cascade, an issue that will be pursued in future work.

PHILIP HALL

Boundary layer stability

The objective of our research is to elucidate basic instability mechanisms (within a theoretical framework) with a view to actively control them.

Work on active control of TS waves, Görtler vortices and crossflow vortices has been completed. We find that TS waves can be made completely stable whilst the other two modes can have their growth rates reduced significantly. An investigation of the nonlinear instability of unsteady pipe flow was carried out. Further work on the wavenumber selection problem for streamwise vortices was completed. Calculations related to the phase equation approach applied to compressible flows were performed. An investigation of the optimal control of boundary layers by suction was begun.

FANG Q. HU

Sound radiation of a turbulent boundary layer

In this work, we aim to better understand noise generation mechanisms in boundary layers. To this end, the impulsively started flat plate serves as a model for the boundary layer in which the parallel assumption is exact.

Thus, we consider sound radiation of a temporally developing turbulent boundary layer. This follows previous work (ICASE report 94-30 by Webb, Otto and Lilley) of non-linear stability of an impulsively started flat plate. We show that the instability waves of the Rayleigh problem provide an "inner" solution for the associated sound radiation problem. By means of matched asymptotic expansions, the far field radiated sound field is found as the "outer" solution in the limit of zero Mach number. Although boundary layer noise has been studied extensively in the literature, past work considered frequency domain analyses and empirical turbulent statistics. The present work is conducted in the time domain and formulated as an initial boundary value problem. The sound field is computed by numerically solving the wave equation in the far field which is in turn reduced to a Klein-Gordon equation for the temporal simulations. This work is conducted in collaboration with M. Y. Hussaini and S. Otto (ICASE).

Sound radiation of mixing layers by vortex methods

Vortex methods offer many attractive advantages for the numerical simulation of flows of practical interest, such as mixing layers and jets. They also have potential in computational aeroacoustics

through their natural connection with the "vortex sound" theory of low Mach number flows. This study explores the viability of applying vortex simulations for computing sound radiation. At present, most vortex methods apply only to incompressible flows. In our approach, the sound radiation from vortex simulations is modeled based on matched asymptotic expansions of the incompressible simulation and compressible far field for small Mach numbers. A model problem of a temporally growing mixing layer has been considered in which the acoustic wave equation is solved in the far field. This work was conducted in collaboration with J. Martin (Christopher Newport Univ.).

Numerical methods for computational aeroacoustics

Our study of numerical schemes for computational aeroacoustics continued with investigations of boundary conditions and boundary schemes (with J. Manthey, graduate student, Old Dominion University). Since the acoustic disturbances are extremely small compared with other fluid variables, special care is necessary at the numerical boundaries where computational accuracy is usually decreased. In particular, numerical implementations of duct walls and non-reflecting boundary conditions are considered in this work.

T. L. JACKSON

Algebraic instabilities in shear flows

Algebraic instabilities arise when some initial disturbances, owing their presence to a finite level of noise present in any flow, grow sufficiently to trigger nonlinear mechanisms or to provide new basic states for secondary instabilities. These instabilities are distinguished from exponential instabilities, where infinitesimal disturbances always grow exponentially in time. The presence of algebraic instabilities may lead to the so-called "bypass mechanisms". Work is continuing on the evolution of disturbances in viscous shear flows. This work offers a means whereby completely arbitrary initial input can be specified and the resulting temporal behavior, including both early time transients and the long time asymptotics, can be determined. The bases for the analysis are: (a) linearization of the governing equations; (b) Fourier decomposition in the plane and streamwise direction of the flow; and (c) direct numerical integration of the resulting partial differential equations. The results provide explicitly both the early time transients and the long time asymptotic behavior of any perturbation. With this knowledge it is then possible to devise means for flow control and it is possible to either delay or enhance disturbances as the need may be.

Two classes of problems have been investigated: (1) family of free shear flows consisting of the jet, wake, and mixing layer; and (2) viscous channel flows. The free shear flows are those that can be examined in an inviscid fluid, thereby allowing some simplification in the system. We found that it was possible to significantly delay the exponential growth in both the jet and wake, but enhancement in the mixing layer was only marginally feasible. This work will appear soon in the *Journal of Fluid Mechanics*. For viscous channel flows, consisting here of the plane Poiseuille and plane Couette flows, many details of the flow field were examined and found important in understanding the dynamics, even at subcritical Reynolds numbers. This work has been submitted for publication.

This method is currently under investigation for the Blasius boundary layer. Future plans will also include viscous pipe flow, as well as the concept of absolute/convective instabilities.

All work was done in collaboration with William Criminale (University of Washington) and D. Glenn Lasseigne (Old Dominion University).

Instabilities in mixing layers

Due to the projected use of the scramjet engine and the observed phenomenon of increased flow stability at high Mach numbers, the investigation of stability characteristics of compressible mixing layers continues to be an active area of research. All previous investigations of the stability of mixing layers, either reacting or non-reacting, have assumed equal molecular weights for the gases above and below the splitter plate. Furthermore, if the flow is assumed to be reactive, all previous analysis employed a simplified kinetic model consisting of a one-step, irreversible Arrhenius reaction. We are currently extending previous results to determine the effects of (i) multi-component gases and (ii) various reduced mechanisms have on the stability characteristics of compressible mixing layers. This work is in collaboration with C.E. Grosch, D.G. Lasseigne and F. Kozusko (Old Dominion University).

M. N. MACROSSAN

Particle Simulation of Near Continuum Flow; Hybrid Equilibrium and Non-Equilibrium Scheme

The Direct Simulation Monte Carlo (DSMC) method is widely used for high speed rarefied gas flows for which the continuum equations of fluid mechanics are expected to be invalid. In many blunt body flows the gas near the windward surfaces is at a high density and the continuum fluid dynamics equations should apply while the flow near the leeward surfaces is at a low density where DSMC should be used. The purpose of this work is to develop a new hybrid low/high density version of DSMC by generalizing the the Equilibrium Particle Simulation Method (EPSM) (D. I. Pullin, *J. Comput. Phys.*, **80**, 231, 1980) to deal with multiple species and varying number of degrees of freedom per particle.

An EPSM subroutine, which can generate a sample of molecular velocities and energies from an equilibrium distribution has been written and has been used to replace the collision routines in an existing DSMC code. Different species of particles are allowed and any number of classical (non-quantum) degrees of freedom.

It has been found that EPSM is about 3 times faster than DSMC for high density flows and thus has great promise for the high density part of a hybrid low density/high density (DSMC/EPSM) code. A description of EPSM and some results of its use are given in an ICASE Interim Report 27. The report contains an analysis of the origin and nature of the dissipation in particle simulation methods when used in the near continuum regime.

The EPSM subroutine will be added to the existing 3-D DSMC code of D. Rault of the GDD-Aerothermodynamics branch at N.A.S.A. LaRC. The hybrid DSMC/EPSM code will be applied to the simulation of the Magellan aero-braking manoeuvre in the atmosphere of Venus in order to understand some aspects of the observed trajectory.

The question of quantum energy of vibration in EPSM needs to be addressed. As a first attempt, this will be done by calculating an effective fractional number of degrees of freedom and including a fraction only of the N excited particles in the re-distribution process.

Development of a Kinetic Theory Based Navier-Stokes Solver

A new finite volume CFD method for the solution of the Navier-Stokes equations, the Equilibrium Interface Method (EIM) (Macrossan & Oliver, *Int. J. Numer. Meth. Fluids*, **17**, 177, 1993) was found in its first order version to be slightly more accurate than a second order Riemann solver based method (Jacobs, *NASA CR-187613*, 1991) and to require approximately half the computational effort (CPU/cell/time-step). However, the indications were that it would require modification to make it more robust, which is the purpose of this work.

The approach taken was to blend the EIM fluxes with those from a related but more dissipative kinetic theory method, the Equilibrium Flux Method (EFM) (Pullin, *J. Comp. Phys.*, **34**, 231, 1980). The blending is based on the second spatial difference of pressure evaluated at cell centers. In smooth regions of the flow, where the pressure gradient is relatively constant the EIM fluxes predominate. At shock waves and other places where the pressure gradient changes rapidly the more dissipative EFM fluxes predominate.

The blended method has been found to be stable for a high Mach number blunt body flows and for the diffraction of a shock wave around a sharp corner where Godunov type methods have been found to exhibit problems (Quirk, *Int. J. Numer. Meth. Fluids.*, **18**, 555, 1994). However, when EIM/EFM is used as the basis of a Navier-Stokes solver a problem of a pressure peak at solid body surfaces has come to light and has not been resolved. A full investigation of the stability of the method is not useful until this problem is better understood.

This problem will be given to a research student at the University of Queensland. Collaboration with D. I. Pullin and J. M. Moschetta at Caltech, who are also working on improved kinetic theory methods to overcome the surface pressure problem (*personal communication*) will be continued.

JAMES E. MARTIN

Three-dimensional vorticity dynamics in a swirling jet model

In spite of their practical importance in combustion applications and jet engines, the nonlinear evolution of swirling jets remains largely unexplored. In addition, the swirling behavior of aircraft wakes recently has become a focal point in investigations of aircraft accidents. Consequently, there is a strong incentive to achieve a better understanding of the fundamental mechanisms governing the behavior of swirling jets and wakes. To this end, we have developed a relatively simple swirling jet model, which nevertheless allows us to study many of the governing dynamical effects analytically and computationally. The model, which represents an extension of the one studied by Caflisch *et al.* (1994), consists of a centerline vortex surrounded by a jet shear layer that contains both azimuthal and streamwise circulation. Linear stability calculations (Martin and Meiburg 1994a) elucidate the growth mechanisms for Kelvin-Helmholtz-type shear instabilities and centrifugal instabilities. Preliminary axisymmetric vortex dynamic simulations based on the Biot-Savart law and Helmholtz's theorem (Martin and Meiburg 1994b) demonstrate the interaction between the two instabilities, which leads to the emergence of counterrotating vortex rings whose circulations are time dependent. Within the present study, we have derived more detailed results for the axisymmetric flow, as well as some first results for helically and three-dimensionally evolving swirling jets.

A three-dimensional computational vortex dynamics technique is employed, which in its specific implementation is essentially identical to the one used in our earlier investigations of mixing layers,

wakes, and nonswirling jets. We consider a temporally evolving flow field, whose vorticity field is discretized into vortex filaments, each of which is represented by a number of nodes along its centerline. The Biot-Savart law is evaluated assuming an invariant algebraic shape function for the vorticity distribution around the filament centerline. One very interesting finding that needs further investigation is the local reversal of the azimuthal vorticity component after a finite time, which results in the formation of a local recirculation zone, in a way similar to the classical observations of vortex breakdown. This result raises some questions about the relationship between the instabilities and the nonlinear evolution of swirling jets and the vortex breakdown phenomenon.

In the future, we will attempt to relate our findings to the phenomenon of vortex breakdown. In addition, the mechanisms of aeroacoustic sound generation will be explored.

ANTHONY T. PATERA

Bayesian–validated computer–simulation surrogates for optimization and design

Although the advent of fast and inexpensive parallel computers has rendered numerous previously intractable calculations feasible, many numerical simulations remain too resource–intensive to be directly inserted into engineering optimization efforts. An attractive alternative to "direct insertion" considers models for computational systems: the expensive simulation is evoked only to construct and validate a simplified input-output model; this simplified input–output model then serves as a simulation surrogate in subsequent engineering optimization studies. As compared to the direct–insertion approach, surrogates offer more complete, efficient, and robust optimization, greater accommodation of prior information, broad applicability to families of objective functions, and a more interactive, flexible design environment.

We present here a "Bayesian–validated" surrogate methodology which permits economical and reliable integration of large-scale (possibly noisy) simulations into engineering design and optimization studies. The distinguishing features of the approach are: an elemental, sequential statistical train–test resampling procedure which permits efficient adaptive construction and rigorous "probably–approximately–correct" validation of proposed surrogates; and a validation–based nonparametric a posteriori error framework which quantifies the effect of surrogate–for–simulation substitution on system predictability and optimality.

The technique has been applied to several heat transfer and fluid flow problems: the selection of the inclusion concentration of a random fibrous composite to achieve a desired effective thermal conductivity, the placement of cylindrical eddy–promoters in laminar–flow heat exchangers to optimize transport, and the design of the shape of axisymmetric bodies to achieve a specified heat transfer rate or (currently, Stokes) drag force. Future work will address extensions both to the algorithmic framework and to the range of applications.

J. R. RISTORCELLI

Modeling issues involving the Reynolds stress and $k - \varepsilon$ turbulence models in flows in which the compressibility of the turbulence is important are being studied. These include flow situations of application to supersonic flight as well as subsonic flight undergoing large accelerations (high lift conditions). The primary concern is with the creation of a consistent set of models for the effects of compressibility. These effects include variable inertia effects associated with variations in the mean density — the mass flux, effects due to the non-zero fluctuating dilatation — the dilatational covariances, dilatational effects due to non-isentropic processes, and fluctuations in fluid properties due to temperature fluctuations. The focus of the research has been on the dilatational covariances which represent a portion of the effects of compressibility on the energy level of the turbulence. Additional work representing the effects of compressibility on the anisotropy and the spectral cascade rate are necessary.

The dilatational covariances

In this last period the focus of investigations in compressible turbulence has been on the effects of compressibility as manifested through the fluctuating dilatation, $d = u_{k,k}$. In particular, we studied the pressure-dilatation covariance, $< pd >$, which exchanges energy between the fluctuating pressure and velocity fields, and the dilatational dissipation covariance, $\frac{4}{3}\nu < dd >$, which acts as an additional source of dissipation of the turbulence energy.

A substantial portion of the time has been spent developing a theory using ideas from acoustics, singular perturbation theory and statistical fluid mechanics to produce explicit expressions for these quantities. The theoretical aspects of this work are now complete. The results of the analysis have been tested in the self-preserving mixing layer; the severe decrease in mixing layer growth as indicated by the well known (if only suggestive) "Langley curve" was found.

Future plans in this area are primarily related to the investigation of additional implications of the analytical work. Codes to test the model and compare it in flows such as the jet — in which the pressure dilatation suppresses turbulence energy, and the wake in which the pressure dilatation enhances the kinetic energy are under consideration. Inspection of the model indicates that it will be able to catch both these behaviors as well as the different behavior seen in the mixing layer versus the equilibrium wall bounded flow. Numerical validation of the theory using data bases of homogeneous compressible flows is to be accomplished. This aspect of the work is complicated by the lack of the existence of a suitable closure for the Reynolds stresses in compressible flows. More importantly, the existence of compressible homogeneous shear data bases with initial conditions adequately reflecting compressible effects due to the vortical motions of the flow, rather than arbitrary initial conditions (as has been the practice) are required.

Non-isentropic sources of compressibility

A nontrivial extension of the theory for the dilatational covariances mentioned above, is to flows that have substantial heat transfer effects. The work has been suggested by the poorly understood effects of heat transfer on the spread spread rates of hot jets. These issues, of application to the mixing problem of jet plumes, have been reported in some of the investigations of the effects of temperature on the spread rate of heated jets by J. Seiner (NASA).

Preliminary results show the importance of the pressure flux and the mean pressure gradient on the sign and magnitude of the nonisentropic corrections to the dilatational covariances.

Progress in this area will be related to the resolution of several difficult issues associated with models for the pressure flux in compressible flows.

Eddy viscosity transport models in compressible flows

The corrections to the eddy viscosity transport expressions used to model turbulent transport of $\{u_i u_j\}, k$ and ϵ are under investigation. This work is suggested by results of compressible wake calculations by J. Morrison (AS&M) and J. Cimbala (Penn State). The issue here appears to be that in compressible flows with large mean density gradients the usual eddy viscosity/gradient transfer models for the turbulent transport in both the $k - \epsilon$ and second-order closure methods appear to be inadequate.

A direct investigation and order of magnitude analysis of various higher order Favre-moment equations for compressible flows with large density variations has shown this to be the case. An analysis carefully accounting for the distinction between Favre and Reynolds averaged quantities in the moment equations has identified the problem.

The work has been suggested by J. Morrison (AS&M) and will supplemented by calculations in two simple flows. Additional work in this area will be conducted in parallel with relevant numerical and experimental investigations with J. Cimbala (Penn State) and J. Bonnet (Poitiers, France) organized by T.B. Gatski (NASA).

The round jet / plane jet anomaly

The inability of a turbulence model to predict the spread rates of both the round and plane jet has never been satisfactorily understood or fully resolved mathematically. A proper turbulence model, constructed in a tensorially consistent way, should not have these problems. The objective of our work is to resolve this anomaly, or at least ascertain that it is not due to a tensorially inconsistent formulation.

In work done with B. Younis (City College, London) it appears that a resolution of the round jet-plane jet spread rate anomaly has been explained. Physical intuition had suggested that a tensorially correct turbulence transport model may resolve this issue. Reynolds stress computations with a more fundamental model for the turbulent transport appear to have verified this conjecture.

This is work currently in progress the details of which will be more fully developed during this next semi-annual period.

Energy transfer in rotating turbulence

The effects of rotation in substantially modifying the spectral cascade of energy are well known. These effects have never been mathematically understood to the point that a successful parameterization of the effects of rotation have been incorporated in Reynolds stress closure setting.

In work done with A. Mahalov (ASU) and Y. Zhou (ICASE) an EDQNM analysis of the Poincare transformed Navier-Stokes equations is underway.

Subsequent to the EDQMN analysis it is expected that suitable arguments about the triadic interactions will lead to the scaling analyses sought. From which point the dependence of one cascade on rotation may be able to be parameterized in a way suitable for Reynolds stress closures.

26

ROBERT RUBINSTEIN

Inertial range scaling of source terms in the theory of weakly compressible turbulence

In the theory of weakly compressible turbulence, the convective derivative of the incompressible pressure field can appear either as a source term for compressible fluctuations (Zank and Matthaeus, 1982) or directly as equal to a compressible field quantity (R. Ristorcelli, 1995). In either case, the development of a model of weakly compressible turbulence must begin with finding the inertial range scaling of the two point correlation of the convective derivative of pressure: in the theory of Zank and Matthaeus, it defines the force correlation from which the compressible damping (the effective viscosity in a model) can be found by renormalization group methods; in the simpler theory of Ristorcelli, the corresponding single point quantity obtained by integration over wavenumbers, will enter a compressible turbulence model directly.

The required correlation is sixth order in the velocity field. Determining the scaling of such a correlation is not trivial; naive power counting based on Kolmogorov scaling can lead to incorrect results: the simplest example (Van Atta, 1972) is the correlation of fluctuating kinetic energy for which power counting predicts a $-\frac{7}{3}$ inertial range scaling instead of the observed $-\frac{5}{3}$ scaling. Analysis of this correlation is now in progress; following Kraichnan and Zakharov, an infinite Kolmogorov spectrum is substituted in the quasinormal approximation common to analytical theories, and the infrared and ultraviolet divergences of the resulting integrals are sought. Infrared divergence implies dependence of the spectrum on the rms velocity, ultraviolet divergence implies dependence on the Kolmogorov scale and hence on turbulent Reynolds number, and convergence is equivalent to the validity of power counting.

Time dependent turbulence modeling

A time dependent turbulence model has been derived from D. C. Leslie's perturbative treatment of the direct interaction theory of shear turbulence. It has the form of coupled relaxation equations for a set of tensors which appear in a decomposition of the Reynolds stress into components with different time dependence. The distortion of wavevector space by the mean shear plays a crucial role in the theory. The theory accounts naturally for the transition in Reynolds normal stress ratios in simple shear flows from short to long times. Application of the theory to normal stresses in oscillating channel flow is in progress.

CHARLES G. SPEZIALE

Turbulence modeling

Recent work has focused on the development and testing of improved Reynolds stress models for compressible turbulent flows of aerodynamic importance. It has become apparent that existing compressible Reynolds stress models yield poor predictions for the normal Reynolds stress anisotropies in high-speed turbulent shear flows. This can seriously impede the ability to predict compressible turbulent flows with separation or secondary flows where the normal Reynolds stress anisotropies play a pivotal role. An entirely new approach is needed for the systematic development of compressible corrections to the representation for the Reynolds stress that will supplant the variable density extensions of incompressible models that have typically been used after invoking the Morkovin hypothesis.

27

Research has been undertaken with R. Abid (High Technology Corporation, LaRC) on the development of models for the deviatoric part of the pressure-strain correlation that account for compressible effects. In an earlier study, the inability to accurately predict Reynolds stress anisotropies in high-speed turbulent shear flows was traced to deficiencies in the modeling of the deviatoric part of the pressure-strain correlation which controls the level of anisotropy. Unlike all previous approaches we have sought to develop algebraic models for the pressure-strain correlation that are systematically obtained from the transport equation for the pressure-strain correlation via the same type of equilibrium hypothesis by which algebraic stress models are derived. The primary point that distinguishes compressible flows from incompressible flows is the fact that the pressure is the solution of a transport equation rather than a Poisson equation. This type of approach naturally brings in a dependence on the turbulence Mach number that is absent in the variable density extension of models.

Both compressible full second-order closures and two-equation models (the latter via the standard algebraic stress approximation) will be developed in the future using this approach. DNS data bases will be used to test the models which will ultimately be applied to practical compressible turbulent flows of aerodynamic interest. The extension of these ideas to the development of subgrid scale models for large-eddy simulations is also envisioned.

SIVA THANGAM

Modeling and analysis of incompressible and compressible turbulent flows

The present study addresses the development and validation of efficient turbulence models for incompressible and compressible flows. Since the limitations in computer capacity and speed often preclude the direct simulation of complex flows of relevance to technical applications, the current practice for high Reynolds number flows of practical interest involves some type of modeling for the turbulent stresses. In this context, the development of two-equation turbulence models that have wide range of predictive capability for separated flows was undertaken from both the analytical and the computational point of view. This is a collaborative work involving Drs. Ye Zhou, R. Ristorcelli and M.Y. Hussaini of ICASE, and Dr. T.B. Gatski of NASA. The work performed to date in this area has established the efficacy of two-equation turbulence models based on the recursion renormalization group theory for the accurate prediction of turbulent flows of practical importance.

The work currently in progress includes analysis of wake flows using the recursion renormalization group theory (r-RNG) based anisotropic two-equation models for incompressible turbulent flows. The model itself was developed during the earlier phase and validated for the benchmark test case of separated flow past a backward-facing step. The work currently in progress has led to the development of recursion-RNG based near wall representation and flow parameter-dependent model coefficients. At present, the model (which includes the near wall effects as well as flow parameter-dependent coefficients for the higher order terms) is being applied to a wide range of external and internal flows and is being validated.

In addition, preliminary investigations on the development of pressure-dilatation and dilatational-dissipation for kinetic energy in the equations of motion for compressible turbulent flows have been

developed and applied to wall-bounded flows and supersonic mixing layers. This effort involving modeling for compressible turbulent flows is being continued.

LU TING

Structural/acoustic interaction problem

The structural/acoustic interaction problem is essential for the prediction and control of panel fatigue and the transmission of external acoustic waves through panels of an airframe into the interior. Experimental investigations were carried out by Dr. Lucio Maestrello at NASA LaRC for the cases where the medium is at rest or moving relative to the airframe. It was found that the acoustic field for a medium moving at constant speed is different from that with acceleration or deceleration. Theoretical studies have been carried out for the interaction problem where the medium is moving at a constant velocity, either subsonic or supersonic. For a medium moving at an unsteady unidirectional velocity $U(t)$, the acoustic field obeys a convective wave equation with variable coefficients depending on $U(t)$, the explicit solution is not available and a comprehensive study of the acoustic field was not made. We are addressing this problem by carrying out a systematic study of the implicit solution of wave propagation in a medium moving at an unsteady velocity.

In collaboration with Drs. F. Bauer of Courant Institute, New York University and L. Maestrello of NASA we are carrying out our theoretical investigation on the interaction problem when the airframe is moving at an unsteady velocity. This problem is much more complex because the explicit solution for the wave propagation can only be obtained in the coordinate system $\bar{x}$ with the medium at rest while numerical solution of the panel oscillation will be much simpler in the coordinate system x moving with the panel. We consider first an initial value problem for an acoustic field induced by an unsteady source distribution, $\bar{q}(t, \bar{x})$ with $\bar{q} \equiv 0$ for $t \leq 0$, in the coordinate system $\bar{x}$ with the medium at rest. We then obtain the implicit solution in the coordinate system x, where the medium is moving with a uniform unsteady velocity $U(t)\hat{\imath}$. Note that signals issued from a point S in the domain of dependence $\mathcal{D}$ of an observation point P, can arrive at point P at time t more than once corresponding to different retarded times, $\tau \in [0, t]$. The number of times of arrival is called the multiplicity of the point S. We show that the multiplicity equals 1 if the velocity U remains subsonic and can be greater when U becomes supersonic. For an unsteady uniform flow $U(t)\hat{\imath}$, rules are formulated for defining the smallest number of I subdomains V_i of $\mathcal{D}$ with the union of V_i equal to $\mathcal{D}$. Each subdomain has multiplicity 1 and a formula for the corresponding retarded time. The number of subdomains V_i with nonempty intersection is the multiplicity m of the intersection. The multiplicity is at most I. Examples demonstrating these rules are presented for media moving at accelerating and/or decelerating supersonic speed. These studies, rules, and examples are presented in a forthcoming ICASE report entitled "Acoustic Field in Unsteady Moving Media" by F. Bauer, L. Maestrello and L. Ting.

Our next step is to develop a numerical program for a given unsteady velocity $U(t)$ to define the domains of integration V_j and compute the retarded time τ_j for a point in a V_j. The next step is to apply our program to compute the unsteady acoustic field for a given source distribution, which can be a simple model or data from experiments. Thus we can compare the acoustic fields induced

by the same source distribution moving with different velocity functions $U(t)$, to demonstrate the effect of unsteadiness and compare with experimental measurements.

NICOLAS VERHAAGEN

Vortical Flows over Double-Delta Wings

The objective of the research is to provide an experimental database for the purpose of understanding the behavior of the vortical flow over double-delta wings and for validating Computational Fluid Dynamics (CFD) methods. The data will help to improve the aerodynamic characteristics of wings to be employed for high-speed and highly-maneuverable aircraft.

Data from an experimental research conducted on a 76/64-deg double-delta wing at the Basic Aerodynamics Research Tunnel (BART) of NASA LaRC has been processed and analyzed. The research was conducted as a cooperation between ICASE, NASA LaRC, the Naval Air Weapons Center (NAWC) and the Delft University of Technology (TUD). Flow visualization tests have provided insight into the effect of angle of attack and Reynolds number on the vortex-dominated flow both on and off the surface of a double-delta wing. Upper surface pressure recordings from pressure tabs and pressure sensitive paint (PSP) have provided data on the pressures induced by the vortices. Flowfield data was obtained at an angle of attack of 10 deg by surveying the flow over the wing using a thin 5-hole probe. The experimental data for this angle of attack was used to validate a numerical solution of the compressible thin-layer Navier-Stokes equations for fully laminar flow computed at NAWC. The solution predicts the flow on and off the surface of the double-delta wing reasonably well. The magnitude of the computed velocities and pressures tends to be lower than that measured in the experiment. Currently, solutions are being developed for the case of fully turbulent flow. To simulate this, the Baldwin-Barth one equation turbulence model is used.

Results of the experimental and numerical research has recently been published and presented at the AIAA 33rd Aerospace Sciences Meeting, held at Reno, Nevada in January 1995. In addition, this study is to be published as ICASE Report No. 95-5.

For a more complete validation of numerical solutions, tunnel wall and model support interference effects should be investigated in more detail both in the experiment and in the numerical simulation. A cooperative research program will be initiated to study these effects.

To further validate CFD solutions, vortex flow surveys at angles of attack other than 10 deg as well as balance measurements will be needed. No schedule has been set yet for these follow-up testings.

ROBERT V. WILSON

Simulation of complex, three-dimensional turbulent jets

Three-dimensional, turbulent jets issuing from elliptic or rectangular nozzles exhibit many complex phenomena including strong azimuthal instabilities, switching of major and minor axes, and increased entrainment rates leading to increased mixing. The objective of the present work is to perform numerical simulations of these flows in order to understand such phenomena.

The study will include Large Eddy and Direct Numerical Simulations of rectangular jets issuing from nozzles and orifices and circular jets with vorticity generating tabs. Past work has focused on developing a code to numerically integrate the governing equations efficiently using high-order finite differencing.

Upon validation of the code, simulations of the above mentioned jets will be performed and compared with the existing numerical and experimental studies.

YE ZHOU

Phenomenological treatment of rotating turbulence

Turbulent flows that are subject to strong rotation are found in many engineering situations, such as turbomachines and the wing tip vortex. Unfortunately, most Reynolds average closure models account inadequately for the effects of strong rotation.

Experimental studies, direct numerical simulations (DNS), large-eddy simulations (LES) and closure approximation have established that solid body rotation suppresses the nonlinear energy cascade from large to small scales. Specifically, uniform rotation causes plane waves to propagate with phase speed $2\Omega k_z/k$. Here the rotation vector is considered to act along the vertical ($\frac{z}{|z|}$). The effect of rotation on the spectrum is through phase scrambling.

In this work, we first note the strong similarity between the magnetohydrodynamic (MHD) turbulence and isotropic turbulence subject to solid body rotation. We then apply the MHD phenomenologies of Kraichnan (*Phys. Fluids* **8**, 1385, (1965)) and Matthaeus and Zhou (*Phys. Fluids* **B1**, 1929, (1989)) to rotating turbulence. When the turbulence is subject to strong rotation, we found that the energy spectrum is $E(k) = C_\Omega(\Omega\epsilon)^{1/2}k^{-2}$, where Ω is the rotation rate, k is the wavenumber, and ϵ is the dissipation rate. The constant C_Ω is found to be related to the Kolmogorov constant and is estimated in the range $1.22 - 1.87$ for the typical values of the latter constant. Our hypothesis on the triple correlation decay rate leads to a spectral law which varies between the '$-5/3$' (without rotation) and '-2' laws (with strong rotation). For intermediate rotation rates, the spectrum varies according to the value of a dimensionless parameter that measures the strength of the rotation wavenumber $k_\Omega = (\Omega^3/\epsilon)^{1/2}$ relative to the wavenumber k. We deduced a 'rule' that relates spectral transfer times to the eddy turnover time and the time scale for decay of the triple correlations. The eddy viscosity is derived with an explicitly dependence on the rotation rate.

A lack of understanding of the fundamental energy transfer process is responsible for the failure of turbulence models in rotating flows. We plan to tackle this problem in the future. Advanced closure theories and numerical simulations will be used to generate the required flow fields.

DAVID C. BANKS

Visualizing vortices in an unsteady flow

Direct Numerical Simulation (DNS) of an unsteady flow can produce hundreds of gigabytes of data. We wish to extract the important flow features from the data and compress their representation enough to fit in workstation memory. This enables the scientist to interactively examine his time-varying data using a local machine. This work is being done in collaboration with Bart Singer.

We used a modified vortex-line to extract vortex skeletons from an unsteady shear flow, calculated over 231 time steps. By representing the cross-sections of the tubes as truncated Fourier series, we were able to compress the representation of the flow substantially.

We plan to enhance the surface representation of the vortex tubes by using dynamic texture-maps and dynamic displacement-maps along the tubes. These techniques require high-end graphics hardware and/or supercomputing capability.

Interactive, immersive visualization of time-varying data

It can be difficult to comprehend the development of 3D data. The problem is made worse by viewing the data on a 2D computer screen, controlling it by a 2D mouse. We wish to give the scientist an immersive view of his 3D data by using stereoscopic display, 3D tracking, and real-time scene updates.

We have developed an interactive, immersive system to view time-varying data. The hardware tools include SGI graphics workstations with texture boards, stereo glasses, and 3D head and hand tracking. The software tools include libraries from the Electronic Visualization Laboratory at the University of Illinois at Chicago, the University of Alberta, and the San Diego Supercomputer Center. The system has been used to investigate the onset of turbulence in a shear flow over a flat plate. It was also demonstrated at the ACM SIGGRAPH conference in Orlando.

We plan to use the system to examine other unsteady flows and to investigate the geometry of a fluttering airfoil. This work has been done in collaboration with Michael Kelley.

GIANFRANCO CIARDO

Efficient description and solution of discrete-state stochastic processes

Discrete-state stochastic processes are useful to model the performance and reliability of a wide class of systems. Formalisms such as stochastic Petri nets (SPNs) or queueing networks (QNs) are used to describe the system at a high-level. Depending on the structure of the state space and on the probability distributions used for the durations of the events modeled, the underlying stochastic process described by the high-level model can be a continuous or discrete time Markov chain (CTMC or DTMC), an independent or ordinary semi-Markov process (ISMP or SMP), a Markov regenerative process (MRGP), or, in full generality, a generalized semi-Markov process

(GSMP). Our goal is to define a flexible and hierarchical language for the description and solution of these models, which (1) enforces strict typing, (2) ensures semantically correct operations on constants, random variables, and stochastic processes, (3) allows an analyzer program to use the more appropriate solution method according to the type of underlying stochastic process and of the measures requested (steady-state vs. transient).

We have completed a first draft defining the SMART (Stochastic Markovian Availability, Reliability, and Timing) Language. This is a declarative language containing features especially targeted to conduct the type of modeling experiments often needed in real applications. For example, it is possible to define arrays of measures, corresponding to the solution of a model for a given set of input values, thus resulting in the definition of a parametric study. Furthermore, if the required solutions do not depend on each other, it is possible to distribute them over multiple workstations, thus achieving a reduction of the wall time required to conduct a modeling study. A unique feature of the SMART Language is the ability to specify explicit fixed-point iterations for the approximate solution of models which have been decomposed in quasi-independent submodels. The implementation of a SMART Analyzer is ongoing at William and Mary, concurrently with this effort. We have a basic prototype running without solution algorithms, which we have been using to test the appropriateness of the syntactic and semantic choices made when defining the SMART Language.

Further work includes two main efforts. First, the definition and implementation of the solvers continues and, as it progresses, we will periodically update the SMART Language to reflect what we learn from its use. Second, we plan to use the resulting package to analyze performance and reliability problems of interest. In particular, we intend to explore the use of continuous and discrete time distributions in practical models, and their effect on the quality of the computed measures.

THOMAS W. CROCKETT

Visualization and graphics for parallel applications

Applications which run on large-scale parallel supercomputers often generate massive output datasets, especially when time-dependent simulations are involved. Visualization techniques are imperative for obtaining a complete understanding of the behavior of many of these simulations. However, the sheer size of the data poses a number of impediments to the traditional approach to visualization, which relies on post-processing with workstation-class systems. Our objective is to develop algorithms and methodologies which can exploit the power of the parallel systems to perform graphics and visualization in-place. This approach is capable of producing live visual output at runtime, thus reducing the need to move large datasets across the network, and providing a visual record of intermediate calculations which would otherwise be discarded because of storage and I/O limitations.

Our work is focused on the development of efficient, scalable parallel rendering algorithms which can be incorporated into graphics libraries for embedded use within parallel scientific and engineering applications. As part of this research, we have developed a prototype library known as PGL, which serves as a testbed for our algorithmic developments. During the current reporting period, the PGL software was ported to the IBM SP2, where it has achieved rendering rates in excess of one million triangles/second using 128 processors. Preliminary results indicate that performance on the SP2 scales somewhat better than expected, but we have also found significant

performance variability, apparently related to the SP2's hardware and software architecture. To better understand these results and their implications, we have begun a detailed experimental investigation of performance on both the SP2 and the Intel Paragon. It is expected that this work will lead to improved algorithms which are more efficient with large numbers of processors, and which have better load balancing characteristics.

We have also developed a new algorithm for retrieving image segments from individual processors' memories and assembling them for efficient network transmission to remote workstations for display. This algorithm provides greatly improved performance with large numbers of processors, reducing the overheads which are inherent in the serialization of a parallel data stream. The new display algorithm and SP2 support were incorporated into an updated version of PGL, which was released to several "beta" users around the country under the auspices of the HPCC program. Performance tuning of the underlying sequential rendering algorithms resulted in improvements ranging from 20-50% for several test scenes. Several applications which incorporated the PGL technology were demonstrated live at the Supercomputing '94 conference. We have also begun a collaborative effort with Rich Greco of Intel's Scalable Systems Division to develop parallel rendering algorithms for future-generation systems, and to enhance the functionality and performance of PGL on current platforms.

Our future work will proceed on several fronts, including algorithms, software, and applications. The near-term focus is on scalability and load-balancing considerations for large numbers of processors (several hundred). We will also be working with David Banks (ICASE) and Will Bene (Old Dominion University and Computer Sciences Corporation) to apply PGL to the visualization of vortex development in transitional flows.

PHILLIP M. DICKENS

Direct execution simulation of parallel message-passing codes

As massively parallel computers become increasingly available, user's interest in the scalability of their code is growing. However, such computers are in high demand, and access to them is restricted. We have developed a system called LAPSE (Large Application Parallel Simulation Environment) which allows a programmer to use a small number of parallel processors to simulate the behavior of a message-passing code running on a large number of processors of a "target" machine. The approach taken in LAPSE is to actually run the application using a separate Unix process for each simulated physical processor. LAPSE provides reasonable slowdowns relative to the natively executing application code. However, as the number of Unix processes per physical processor increases, the cost of context switching and message passing increases. This can become a performance bottleneck if there are many processes per physical processor. Current research, performed in conjunction with Piyush Mehrotra and Matthew Haines, is focusing on improving the performance of LAPSE when there are multiple processes per physical processor.

Our approach is to integrate the LAPSE system with Chant, a threads package being developed at ICASE. Chant allows multiple threads of control on a given physical processor, where the multiple threads are run *collectively* as a single Unix process. Thus the simulated physical processors in LAPSE may be implemented as threads rather than Unix processes, significantly reducing the cost of context switching and message passing. There is one significant problem which we encountered in

this implementation. Each of the LAPSE threads is simulating a physical processor, and thus each thread must have its own private data. Thread packages such as Chant provide a shared memory model only. To solve this problem we are implementing each thread as a C++ class object, where each instantiation of the class object is automatically provided its own private copy of the class data.

Our future plans are to complete our implementation of the LASPE/Chant interface, and to develop a test suite to compare the performance of our new system and the LAPSE system. Also, we will study the issue of automatically converting a LAPSE application code into a C++ code which will provide thread private data.

MATTHEW HAINES

Runtime support for task and data parallelism

Integrating task and data parallelism requires sophisticated runtime support to handle issues of communication and synchronization among parallel tasks in an application, each of which may execute in a data parallel manner. Additionally, mapping these tasks onto a limited set of physical resources may require sharing resources among several tasks, which may or may not be related. Our objective is to design and implement a thread-based runtime system that can provide an efficient solution to the problems of integrating task and data parallelism. This work is being done in collaboration with Piyush Mehrotra, and students David Cronk and Bryan Hess.

We have divided the runtime project into two layers, a lower layer for supporting language-independent, lightweight threads capable of communication in a distributed memory environment (Chant), and a higher layer for providing the support required by the Opus language (Opus Runtime). We have designed and implemented Chant atop POSIX pthreads and MPI, and are currently running on the Intel Paragon and a network of workstations. In addition to supporting Opus Runtime, we are using Chant to determine the benefits of combining threads with data parallel compilers and scientific applications.

We plan to continue development of these two layers of the runtime system, integrating them at some point in the near future. We also plan to use the Chant layer as a vehicle for studying issues related to load balancing, irregular scientific problems, and thread-based performance prediction and evaluation.

Smartfiles: An object-oriented approach to data file interoperability

Data files for scientific and engineering codes typically consist of a series of raw data values whose interpretation is buried in the programs that interact with these files. In this situation, where the meaning of a data file is implicit in associated programs, making even minor changes in the file structure or sharing file between programs (interoperability) can only be done by careful examination of the I/O statements of the programs using this file. Moreover, parallel changes are often required at many points in the system. By applying object-oriented techniques to files, we can add the intelligence required to improve data interoperability and provide an elegant mechanism for supporting complex evolving or multidisciplinary applications. This work is being done in collaboration with Piyush Mehrotra and John Van Rosendale.

We have completed the preliminary design and prototype software system, and are currently working on extending the prototype to support the full functionality of the design.

We plan to continue development on the Smartfile system and study the performance characteristics of the system as compared with existing approaches. We also plan to start using the system for some of the application codes being used at ICASE.

JIM E. JONES

Multilevel applications to porous media calculations

The simulation of flow in porous media has important environmental applications where fluid velocities must be calculated accurately. Because of the geology of the media, irregular grids are often needed for accurate computations. There is a need for an efficient method for solving the very large set of discrete equations used to model the flow.

Extending work begun before arriving at ICASE in March 1995, I have developed a new multilevel approach for the efficient solution of the discrete set of equations that arise from the mixed finite volume element (FVE) discretization of the governing equations for flow in porous media. The approach uses the standard finite difference discretization as a preconditioner for the more accurate mixed FVE discretization. Preliminary testing shows that the new approach has potential as a fast solver for practical simulations.

The multilevel approach, used here for porous media calculations, can be broadly applied. We will be implementing it on the Langley parallel testbeds.

DAVID E. KEYES

Parallel algorithms of Newton-Krylov-Schwarz type

Parallel implicit solution methods are increasingly important in aerodynamics and other fields leading to large nonlinear systems with sparse Jacobians. Several trends contribute to their importance. Multidisciplinary analysis and optimization require rapidly achievable low residual solutions, since individual component codes are often iterated and their results differenced for sensitivities. Problems possessing multiple space or time scales motivate implicit algorithms, and arise frequently in locally adaptive contexts and in dynamical contexts such as aero-elasticity. Meanwhile, the demand for resolution and prompt turnaround forces consideration of parallelism, and, for cost effectiveness, the high-latency, low-bandwidth parallelism available from workstation clusters. Our ICASE research program in Newton-Krylov-Schwarz (NKS) solvers is responding to this need, in collaborations with academia, national laboratories (NASA and DOE), and industry (Boeing and UTRC).

An NKS method combines a Newton-Krylov (NK) method, such as nonlinear GMRES, with a Krylov-Schwarz (KS) method, such as additive Schwarz. The linkage is the Krylov method, whose most important characteristic, from a computational point of view, is that information about the underlying Jacobian needs to be accessed only in the form of matrix-vector products in a relatively small number of directions. However, if the Jacobian is ill-conditioned, the Krylov method will require an unacceptably large number of iterations. The system can be transformed through

the action of a preconditioner whose inverse action approximates that of the Jacobian, but at smaller cost. It is usually in the choice of preconditioning that the battle for low computational cost and scalable parallelism is won or lost. In KS methods, the preconditioning is introduced on a subdomain-by-subdomain basis, providing well load-balanced data locality for parallel implementations over a wide granularity range. A two-grid-level form of additive Schwarz provides a mesh-independent and granularity-independent convergence rate in elliptically dominated problems, including nonsymmetric and indefinite problems.

Encouraging applications of NK technology have arisen in collaborations with W. K. Anderson of the Computational Aerodynamics Branch at NASA Langley. Anderson's state-of-the-art unstructured-grid Euler and Navier-Stokes codes employ a solver common in form to that of many implicit codes. A left-hand side matrix (not a true Jacobian) is created, in whose construction computational short-cuts are employed, and which may be stabilized by a degree of first-order upwinding that would not be acceptable in the discretization of the residual, itself. Inconsistency between the left- and right-hand sides prevents the use of large time steps. Using the built-in solver capability as the preconditioning and matrix-free approximations to the true Jacobian, time steps can be built up to large values, recovering Newton's method in the limit. This work has been extended to three-dimensional problems in conjunction with E. J. Nielsen and R. W. Walters of Virginia Tech. Related projects include a two-level Schwarz preconditioned version of TRANAIR with D. P. Young* and R. G. Melvin of Boeing, a low Mach number combustion simulation based on NKS with D. A. Knoll* and P. R. McHugh of INEL, an unstructured Euler code parallelized on the Paragon, the SP2, and workstation clusters with Venkatakrishnan of ICASE and S. Kareti of ODU, an Ethernet Sparcstation cluster implementation of a KS solver for a structured-grid Euler Jacobians with M. D. Tidriri of ICASE, W. D. Gropp* of Argonne, and S. Balay of ODU, an NKS solver for a model full potential equation parallelized on the SP2 with X.-C. Cai* of UC-Boulder and Gropp, and the Eulerian part of a semi-implicit/semi-Lagrangian formulation of the shallow water equations in ocean modeling with J. Chefter-Levin* and C. K. Chu of Columbia. (Asterisked collaborators have visited ICASE during the past reporting period.)

Thus, a variety of CFD applications are (or have inner) nonlinear elliptically-dominated problems amenable to solution by NKS algorithms, which are characterized by relatively low storage requirements (for an implicit method) and locally concentrated data dependencies. The main disadvantage of NKS algorithms is the large number of parameters that require tuning. Each component (Newton, Krylov, and Schwarz) has its own set of parameters. Parametric tuning is important to performance and will be the focus of the next phase of research.

Communication modeling in distributed computing

The so-called "hyperbolic model" for communication in multi-layer contended networks, introduced by I. Stoica, F. Sultan and myself, has been successfully applied to the multi-stage interconnection network on the IBM SP2. The model represents a compromise between communication models of greater fidelity, which are intractable for general purpose use in performance prediction, and trivial models (like the PRAM), which have no value in predicting the performance of buildable hardware.

The hyperbolic model is a set of rules (a "communication calculus") for building up message service time estimates from parameters characterizing system components and from a given message pattern. Contention is permitted; however, the simultaneous execution of randomly interfering

parallel jobs is excluded. The communication system is modeled as a directed communication graph in which terminal nodes represent the application processes that initiate the sending and receiving of the information and in which internal nodes reflect the layered structure of the underlying communication architecture. The direction of graph edges specifies the flow of the information carried through messages.

The test codes employed in ICASE TR 94-78 (containing both time and space parallelism) have been converted to MPI and are now being run on several additional diverse architectures, including three different cluster environments that support MPI, in order to obtain some hard data for use in evaluating the relative suitability of clusters for NASA's quotidian computing needs.

SCOTT T. LEUTENEGGER

Data base support for subset retrieval and visualization of scientific data

The objective of this research is to design and implement a prototype database to facilitate retrieval of subsets of large scientific data sets. The data subsets are anticipated to be used as inputs to other codes, such as for visualization or MDO.

Currently many scientist store and retrieve data sets as files. When the scientist is interested in a subset of the data they read in the entire data set and strip out the portion of interest. This is not practical when data sets are large. Our approach is to provide database support to retrieve only those pages from disk that contain the desired data. Typical CFD data sets are two or three dimensional, thus we provide a multi-attribute indexing technique.

During a two week visit to ICASE in December 94 we continued development of our prototype by improving the efficiency of filtering FAST format data into our canonical DB format, implementing skip queries, and implementing point queries.

We will continue development of the prototype this summer by first extending the system to handle multi-block grids. We will then begin work on implementation of indexing support for irregular grids. Supporting indexing of irregular grids will be done using R and R* trees and/or grid files using the bulk loading algorithm of Nicol and Leutenegger.

KWAN-LIU MA

Parallel volume ray-tracing unstructured-grid data on distributed memory architectures

Computational modeling of scientific and engineering problems with complex geometries often uses finite volume methods or finite element approximations on unstructured grids. Typically, the problem domain is decomposed into small cells, called elements. Popular element types include the tetrahedron, triangular prism (pentahedron) and hexahedron. Many visualization techniques have been developed for the interrogation and analysis of unstructured-grid data. While exterior face rendering and cutting plane methods remain the most common and affordable techniques, three-dimensional methods such as direct volume rendering have received considerable attention because they can capture the overall data domain in a single image.

However, large-scale simulations can generate data with more than hundreds of thousands of elements. The absence of a simple indexing scheme for three-dimensional unstructured grids makes

direct volume rendering a computationally expensive process. As parallel processing enables the solution of many other compute-intensive problems, computer graphics and visualization researchers have also been exploiting various parallel architectures. In this research, we are developing a distributed parallel volume ray-tracing algorithm for visualizing unstructured-grid data.

Our algorithm differs from previous ones in several ways: it is completely distributed, less view-dependent, highly scalable, and more flexible. First, both the data and the rendering computation are distributed across the available processing nodes. Inter-processor communication is only needed for the image compositing step. At each processor, ray-tracing of local data is performed independent of other processors. Image compositing is overlapped with the ray-tracing processes to achieve higher parallel efficiency. Second, the overhead due to view changes is kept to a minimum since a good distributed rendering algorithm must cope with frequent view changes to support truly interactive data exploration. Third, while using more processing nodes increases the number of image compositing layers, the image area that each processor must handle decreases. As a result, the algorithm is very scalable. Finally, although the prototype implementation handles only tetrahedral cells, the algorithm can be generalized to handle a mix of cell types and arbitrary object geometry, such as objects with holes and concavities.

Preliminary performance studies on the Intel Paragon clearly indicate many opportunities for further optimization of the algorithm and its implementation. First, data partitioning for rendering as a postprocess should be further studied to reduce the current 40% load imbalance. Dynamic load balancing, though more effective, is harder to implement in a distributed environment. We will experiment with static load balancing using interleaving. Next, as we approach interactive rendering rates using more processors or more powerful processors, we may need to reevaluate both the image-space partitioning and image compositing step as the rendering step becomes less dominant. Finally, a performance model should be derived for more thorough studies of this algorithm and its implementation. Other future work includes porting the renderer to clusters and more power machines like the IBM SP2; handling I/O in a more efficient manner; and implementing support for runtime monitoring of parallel simulations.

KURT MALY

pPVM on multiple FDDIs

pPVM is an effort to provide high performance communication support in a cost effective way for parallel and distributed computing. Many applications are capable of pushing data at a rate higher than the bandwidth available to a workstation due to limitations in the physical media of the underlying network. The communication network becomes the bottleneck and the classes of applications that can be effectively solved on the cluster is significantly reduced. We use a cost-effective approach of parallel networking to improve the communication bandwidth of the network and make cluster computing effective for a larger class of applications.

In the past we have achieved a parallel PVM implementation using ethernets and demonstrated substantial reduction in the communication component of a computation. This year our effort concentrated on designing and implementing adaptive scheduling algorithms for multiple FDDI rings. Two major problems which we are now addressing are the sensing of the load on an FDDI ring from the PVM level and the saturation of the ring with limited number of workstations (less

than 10). We have succeeded in modifying pPVM to work on heterogeneous networks (both FDDI and Ethernet) and have preliminary methods for sensing load independent of the particular FDDI card used in the network.

In the near future we will be testing various scheduling algorithms. These algorithms will be tested on a data-transpose kernel. This kernel represents the communication requirements for a large class of real applications such as solution of PDEs, FFT computation, sorting, etc. Next on our schedule is the comparisons of real applications running over PVM vs pPVM on heterogeneous networks using at least two FDDI rings. This research was conducted in collaboration with Mohammed Zubair and Shubhangi Kelkar.

PIYUSH MEHROTRA

Integration of task and data parallelism

High Performance Fortran has targeted data parallel algorithms. However, many scientific applications such as multidisciplinary optimization codes, exhibit a coarser grained task parallelism at the outer level while being data parallel at the inner level. We have designed a set of language extensions, called Opus, for exploiting both levels of parallelism. The objective is to provide a high-level programming environment for implementing such codes in a plug-compatible manner on a network of parallel and distributed machines. This work is being done in collaboration with Matthew Haines, John Van Rosendale, Barbara Chapman and Hans Zima.

The first version of Opus provided coarse-grain task parallel extensions for data parallel languages such as Vienna Fortran and HPF, and was based upon two major concepts: *tasks* as asynchronous activities, and *SDAs*, which represent a pool of common data together with methods for controlled access to these data items. We used these construct to encode the outer levels of a simple multidisciplinary optimization code. This exercise resulted in an overall review and generalization of the language definition. The most significant change was the generalization of the *method* concept to allow *synchronous* as well as *asynchronous* activations along with mechanisms to wait for their completion. Asynchronous SDA method activations provide the functionality of tasks – which therefore are no longer needed as a separate concept. Furthermore, the syntax and semantics of SDA variable declaration and associated object creation were significantly extended, and the role of persistent objects in the language as well as the mechanisms for writing and reading such objects were clarified. We are currently modifying the runtime support system in order to take into account the changes in the language.

We plan to continue evaluating the language design, modifying it as necessary. We are implementing the Opus Runtime system on top of Chant, the thread based portable language independent system that we are developing. We also plan to build a source-to-source translator so that Opus programs can be directly transformed to target the runtime system.

Evaluation of HPF

The stated goal of High Performance Fortran (HPF) was to "address the problems of writing data parallel programs where the distribution of data affects performance." We have been using data parallel codes of interest to NASA to evaluate the effectiveness of the language features of HPF. This work is being done in collaboration with graduate student Kyle Winn.

40

After analysis, we restructured TLNS3D so as to be suitable for HPF. TLNS3D is a multi-block and hence exhibits parallelism at two levels: across blocks and within each block. We used APR's HPF compiler for the Intel Paragon to study the performance of the HPF code. The overall wall clock times were disappointing, remaining constant as we increased the number of problems. The computation is divided into two phases: in the first phase values are computed independently on each block and then boundary values for abutting blocks are exchanged. Analysis of the generated code showed that with increasing number of processors the first phase scales, however the time for the second phases increases. This is because the boundary data exchange is data dependent and hence the compiler cannot make any static optimizations and currently does not have the runtime support to execute this phase efficiently.

We plan to continue our evaluation of TLNS3D and other codes, in particular studying the performance of these HPF codes as the compilers become more sophisticated.

DAVID MIDDLETON

Further extensions to a simple combining network

Communication operations often take a large fraction of the time spent by parallel computations, so it is appropriate to study how communication networks might be made faster, or more powerful. A combining network has been defined that requires minimal hardware resources. An executable specification has been created for this network that demonstrates the powerful processing it can perform within the communication operations. Future extension of this network involves several potential directions.

The current direction relates to the fact that such a network requires dedicated hardware elements to achieve speeds near its potential. I have been studying Field Programmable Gate Arrays and similar programmable logic devices with a view to using them to construct both the network nodes and the interface units that connect the PEs to the network.

For future work, a second direction is to study the kind of message processing to be performed in the PEs as they receive a stream of messages. At the moment, each PE accepts a particular message by its position in the stream. While the processing has to occur at hardware speeds, rather than through a software filter, this still allows for additional possibilities, such as waiting for a specific pattern. These extensions turn out to be quite powerful for some graph matching algorithms (developed by Bruce T. Smith), and so deserve to be added to the specification. A third direction for extending the system is to introduce interrupts to the network. In the context of its intended use, several networks are configured dynamically in a system during execution. To prevent a slow computation from blocking others, its network's operation may occasionally be suspended so it can be moved to a more convenient set of nodes. This imposes constraints on the kind and location of state information that can be meaningfully kept during a message wave (given that the physical network nodes may be reassigned). These interrupts add a new layer of behavior to the system that in not fully modeled at the moment.

41

DAVID NICOL

Utilitarian parallel simulator

One of the principle reasons that parallel discrete event simulation has not made an impact on simulation practice is the difficulty of programming the synchronization correctly. Our objective is to alleviate this problem in the context of simulating parallel computer and communications systems.

Our approach recognizes that several successful conservative synchronization algorithms are known to be effective in the context of simulating parallel computer and communication systems. We have developed a software library called U.P.S. (Utilitarian Parallel Simulator) that is used in conjunction with the commercial CSIM simulator package, to provide transparent synchronization and communication. The simulation modeler develops models largely as he would for a serial simulation. By incorporating U.P.S. constructs at the interface between simulation processes, the synchronization and communication activity is carried out automatically. U.P.S. has been implemented on the Paragon.

Future plans include porting U.P.S. to the IBM SP-2, changing it to use MPI instead of nx, and to use it in the Chant Threads Simulation Project. This research was conducted in collaboration with Phillip Heidelberger (IBM).

Chant threads simulation project

System software that supports threads is difficult to develop because there are numerous design decisions whose impact cannot be forecast without implementation and experimentation. Our objective is to develop a simulation tool that will be used as part of the system software design process, in order to reduce the effort needed to make rationale design decisions.

Our approach has been to develop the Threaded Application Parallel System Simulator (TAPS). It presently supports a simple model of threaded workload under the Chant System, but does provide the capability of investigating different policies with regard to scheduling and polling for messages. TAPS is written in C++ and uses CSIM; its design supports easy extension to other workload and communication network models.

Our plans for the future are to parallelize TAPS using U.P.S., to expand the types of thread scheduling policies it supports, and to develop more sophisticated models of thread workload. This research was conducted in collaboration with Matthew Haines.

Reliability and performance integration

Reliability and performance modeling has hitherto been disjoint activities, leading to disjoint models with results that do not obviously mesh. We address this problem by investigating how reliability and performance information can be extracted from the same modeling tool.

Our approach is to integrate our REST interface reliability engine into the commercial BONeS Designer system simulator. The design engineer designs a system using a design tool to which he is accustomed, and from which he customarily extracts performance data. With the integration of REST, reliability information about this same model is produced automatically.

Future plans are to expand the set of models developed using this tool, and to develop automated techniques for model reduction.

ALEX POTHEN

Ordering algorithms for solving linear systems of equations medskip

The time and storage required to solve large, sparse, systems of linear equations strongly depend on the ordering of the equations and unknowns. Direct solvers require orderings that reduce fill, frontal methods require the reduction of a parameter called the mean wavefront, and preconditioned iterative solvers require the reduction of the envelope size of the associated coefficient matrix. Remarkably, an algebraic method is capable of computing good orderings for all three classes of problems, and we design and evaluate the resulting ordering algorithms.

The algebraic algorithm for reducing the envelope of sparse matrices is as follows: First, we compute a specified eigenvector of the Laplacian matrix associated with the given matrix, sort the components of the eigenvector into non-increasing (or non-decreasing) order, and then permute the matrix according to this ordering. Together with Alan George (Waterloo), we have provided a *raison d'être* for the spectral envelope-reduction algorithm by formulating the problem as a quadratic assignment problem, and then showing that the spectral algorithm finds an approximate solution. This technique shows that the spectral orderings are nearly optimal for the related 2-sum problem. This work has applications in computing incomplete-factorization-preconditioners, for wavefront reduction in frontal methods for structural analysis, and in sequencing problems in computational biology.

An improved hybrid algorithm for envelope and wavefront reduction, which combines the good "global" framework of a spectral ordering with the muscle of a "local" combinatorial ordering has also been developed with Gary Kumfert, a Ph.D. student at Old Dominion University. Our results show that the hybrid algorithm can (1) reduce the envelope parameters significantly over other algorithms, and (2) calculate the ordering faster than a spectral algorithm, since the expensive eigenvector can be computed to lower precision. Dr. Jennifer Scott at the Rutherford Appleton Laboratory (Oxford, England) has provided us with preliminary results from the use of these orderings within a frontal solver, showing that the hybrid orderings can reduce the solution times on a Cray Y-MP by a factor of two to three.

A spectral nested dissection ordering algorithm (SND) has also been designed for fill-reduction. This algorithm is based on a spectral partitioning algorithm that partitions vertices into two sets according to their Laplacian eigenvector components. SND computes good orderings for solving linear systems of equations by direct methods in parallel. Together with Ed Rothberg (Silicon Graphics) and Lie Wang (a former Ph. D. student now at IBM), we have evaluated the performance of this ordering for solving positive definite systems of equations on the Intel Paragon and the SGI Power Challenge. Our results show that SND outperforms other ordering algorithms (minimum degree and other nested dissection algorithms) by a wide margin for many large problems from diverse application areas such as fluid flow, structural analysis, and financial modeling.

Currently we are working on a library-quality software package incorporating these ordering algorithms. The influence of the ordering on the quality of preconditioners, application to anisotropic problems, and partitioning physical problems in some preferred physical direction, are other issues being studied. Future work will focus on computing these orderings in parallel.

A Microeconomic Scheduler for Parallel Computers

We have initiated work on a novel scheduling algorithm based on a microeconomic paradigm for scheduling jobs on a parallel computer. (This is joint work with Ion Stoica and Hussein Abdel-Wahab, both at Old Dominion University.) The microeconomic paradigm has been applied to problems in distributed computing by a group of researchers at Xerox Palo Alto Research Center and M.I.T. in the last few years. Our work differs significantly from this earlier work in that we do not use an auction mechanism at every time-slice to schedule jobs.

Every user has a *savings account* in which he receives "money" at a constant rate. When a user wants to run a job, he creates an *expense account* for that job to which he transfers money from his savings account. The job uses the funds in its expense account to obtain the system resources it needs for execution. The share of the system resources allocated to the user is directly related to the rate at which the user receives money; the rate at which the user transfers money into a job expense account controls the job's performance.

We have proved that starvation (i.e., a job waiting for ever without being scheduled to run) is not possible in our model. Simulation results show that our scheduler improves both system and user performances in comparison with two current scheduling policies, particularly when the work load is high. The scheduler allocates resources fairly among the users (proportional to their resources). Each user can control the relative performance among his jobs by controlling the rate at which money is transferred into a job's expense account.

We are excited about extending the microeconomic scheduler to simultaneously allocate processors, memory, and communication bandwidth to jobs on a parallel computer. We are also developing analytical models to understand the performance of this scheduling algorithm.

MOULAY DRISS TIDRIRI

Newton-Krylov-Schwarz algorithms in CFD

Newton-Krylov methods are potentially well suited for the implicit solution of nonlinear problems whenever it is unreasonable to compute or store a true Jacobian. Krylov-Schwarz iterative methods provide good data locality for the parallel implicit solution of PDEs that arise in CFD. We have continued our investigation of the Newton-Krylov-Schwarz (NKS) algorithm, which is a hybrid of the two. Several calculations of 2D and 3D compressible Euler flows have been performed.

We are also exploring parallel implementation and performance aspects of NKS methods with with W.D. Gropp (Argonne National Laboratory), D.E. Keyes (ICASE & ODU), and L. C. McInnes (Argonne National Laboratory).

In the future we are planning to apply this methodology to the computation of 3D Navier-Stokes flows.

Theory and application of Schwarz/MUSCL schemes

Our goal here is to provide an optimal two-level Schwarz/MUSCL type scheme.

We have established the theory and a convergence analysis of the hybrid Finite Element/Finite Volume schemes, and performed several experiments to verify the theory. Based on the error estimates developed in the above study, we have also proved the convergence of the combination of

Schwarz domain decomposition methods with Finite Volume schemes, and established the rate of its convergence.

We are planning to implement the two-level MUSCL-Schwarz algorithms, which are intended for the parallel implicit solution of aerodynamic problems, and demonstrate their optimality.

WILLEM VERMEER

Three-dimensional shock wave visualization on unstructured grids

Shocks occur when an aircraft travels close to or faster than the speed of sound. Their study is important, both because of their direct effects on the flow field, and because of the problem of sonic booms. The goal of our research is to develop effective algorithms for locating and displaying shock waves in CFD-data, which can be used for any type of compressible flow, and for either vertex-based or cell-centered data on unstructured grids. The algorithm used should be robust and require minimal user intervention in detecting shocks. Then once the shocks are detected, the user should be able to interactively examine their three-dimensional structure and topology relative to the aircraft.

In pursuing this goal, we implemented three different algorithms for shock detection, based on local Mach-number, density gradients, and directional derivatives, respectively. Our first approach was to use isosurfaces of the local Mach-number. This works, but can give erroneous results with oblique shocks. The second method looks for the parts of the flow field having the highest density gradients. Our third method is to look for the Mach 1 isosurface of the flow component in the direction of the density gradient. The latter two approaches work well some of the time, though their performance is sensitive to the nature of the dataset used. For transonic flows having weak shocks, density gradients provide insufficient information to properly locate the shocks. There is also a difficulty in that the difference stencils being used are fairly large. We are thus looking at a scheme for second derivative computations that uses a smaller stencil, thereby reducing the smoothing effect of a large stencil.

Once the shock has been detected, it needs to be visualized. The detected shock will generally be represented as a triangulated surface, and there are then any number of visualization techniques that can be applied to these triangles. Currently, the shock-triangles with the aircraft are hardware Gouraud shaded. An alternative to representing the shock would be to do direct volume rendering, which may give a more accurate representation of the numerical shock. This work is being done in collaboration with Kwan-Liu Ma and John Van Rosendale

LINDA WILSON

Automated load balancing of parallel discrete event simulations

A key problem in parallelization of discrete event simulations is the fact that workload is irregular and hard to predict. We are addressing this problem by developing automated load balancing methods for the SPEEDES parallel simulator.

Our accomplishments to date include porting the SPEEDES simulator to the Intel Paragon, and bringing up several large SPEEDES applications, including one developed by Mitre. We have

tuned SPEEDES performance on the Paragon and studied the issue of rapid GVT calculation. We have also begun to study the effect of different load balancing algorithms on the performance of the Mitre application.

We plan to continue studying load balancing algorithms, and migrate the most successful ones into the SPEEDES tool. This research was conducted in collaboration with David Nicol.

REPORTS AND ABSTRACTS

George, Alan, and Alex Pothen: *An analysis of spectral envelope-reduction via quadratic assignment problems*. ICASE Report No. 94-81, November 16, 1994, 25 pages. Submitted to SIAM Journal on Matrix Analysis and Its Applications, September 1994.

A new spectral algorithm for reordering a sparse symmetric matrix to reduce its envelope size was described in [2]. The ordering is computed by associating a Laplacian matrix with the given matrix and then sorting the components of a specified eigenvector of the Laplacian. In this paper, we provide an analysis of the spectral envelope reduction algorithm. We described related 1- and 2-sum problems; the former is related to the envelope size, while the latter is related to an upper bound on the work involved in an envelope Cholesky factorization scheme. We formulate the latter two problems as quadratic assignment problems, and then study the 2-sum problem in more detail. We obtain lower bounds on the 2-sum by considering a projected quadratic assignment problem, and then show that finding a permutation matrix closest to an orthogonal matrix attaining one of the lower bounds justifies the spectral envelope reduction algorithm. The lower bound on the 2-sum is seen to be tight for reasonably "uniform" finite element meshes. We also obtain asymptotically tight lower bounds for the envelope size for certain classes of meshes.

Costiner, Sorin, and Shlomo Ta'asan: *The algebraic multigrid projection for eigenvalue problems; backrotations & multigrid fixed points*. ICASE Report No. 94-82, October 13, 1994, 18 pages. To be submitted to Applied Mathematics Letters, 1994.

The proofs of the theorem for the algebraic multigrid projection (MGP) for eigenvalue problems, and of the multigrid fixed point theorem for multigrid cycles combining MGP with backrotations, are presented. The MGP and the backrotations are central eigenvector separation techniques for multigrid eigenvalue algorithms. They allow computation on coarse levels of eigenvalues of a given eigenvalue problem, and are efficient tools in the computation of eigenvectors.

Otto, S.R., T.L. Jackson, and F.Q. Hu: *On the evolution of centrifugal instabilities within curved incompressible mixing layers*. ICASE Report No. 94-83, October 13, 1994, 30 pages. Submitted to Journal of Fluid Mechanics.

It is known that certain configurations which possess curvature are prone to a class of instabilities which their 'flat' counterparts will not support. The main thrust of the study of these centrifugal instabilities has concentrated on curved solid boundaries and their effect on the fluid motion. In this article attention is shifted towards a fluid–fluid interface observed within a curved incompressible mixing layer. Experimental evidence is available to support the conjecture that this situation may be subject to centrifugal instabilities. The evolution of modes with wavelengths comparable with

the layer's thickness is considered and the high Taylor/Görtler number régime is also discussed which characterizes the ultimate fate of the modes.

Sidilkover, David: *A Genuinely multidimensional upwind scheme and efficient multigrid solver for the compressible Euler equation.* ICASE Report No. 94-84, December 7, 1994, 36 pages. To be submitted to the Journal of Computational Physics.

We present a new approach towards the construction of a genuinely multidimensional high-resolution scheme for computing steady-state solutions of the Euler equations of gas dynamics. The unique advantage of this approach is that the Gauss-Seidel relaxation is stable when applied directly to the high-resolution discrete equations, thus allowing us to construct a very efficient and simple multigrid steady-state solver. This is the only high-resolution scheme known to us that has this property. The two-dimensional scheme is presented in detail. It is formulated on triangular (structured and unstructured) meshes and can be interpreted as a genuinely two-dimensional extension of the Roe scheme. The quality of the solutions obtained using this scheme and the performance of the multigrid algorithm are illustrated by the numerical experiments. Construction of the three-dimensional scheme is outlined briefly as well.

Blackaby, Nicholas D., and Philip Hall: *The nonlinear evolution of the inviscid secondary instability of streamwise vortex structures.* ICASE Report No. 94-85, November 16, 1994, 25 pages. Submitted to Phil. Trans. of the Roy. Soc. of London.

The weakly nonlinear evolution of an inviscid marginally unstable wave growing on a boundary layer supporting a streamwise vortex structure is investigated. The nonlinear growth of the wave is found to be controlled by the diffusion layer located at the edge of the critical layer associated with the wave. The evolution equation is found to depend on the upstream history of the wave and the solution of the equation suggests that the wave either restructures the mean state so as to make it stable or develops a singularity at a finite distance downstream of the point of neutral stability.

Coward, Adrian V., and Philip Hall: *The stability of two-phase flow over a swept-wing.* ICASE Repo No. 94-86, November 16, 1994, 43 pages. Submitted to the Journal of Fluid Mechanics.

We use numerical and asymptotic techniques to study the stability of a two-phase air/water flow above a flat porous plate. This flow is a model of the boundary layer which forms on a yawed cylinder and can be used as a useful approximation to the air flow over swept wings during heavy rainfall. We show that the interface between the water and air layers can significantly destabilize the flow, leading to traveling wave disturbances which move along the attachment line. This instability occurs for lower Reynolds numbers than is the case in the absence of a water layer.

We also investigate the instability of inviscid stationary modes. We calculate the effective wavenumber and orientation of the stationary disturbance when the fluids have identical physical

properties. Using perturbation methods we obtain corrections due to a small stratification in viscosity, thus quantifying the interfacial effects. Our analytical results are in agreement with the numerical solution which we obtain for arbitrary fluid properties.

Carr, Eric, and David Nicol: *Empirical study of parallel LRU simulation algorithms.* ICASE Report No. 94-87, November 16, 1994, 15 pages. Submitted to 1995 Workshop on Parallel and Distributed Simulation.

This paper reports on the performance of five parallel algorithms for simulating a fully associative cache operating under the LRU (Least-Recently-Used) replacement policy. Three of the algorithms are SIMD, and are implemented on the MasPar MP-2 architecture. Two other algorithms are parallelizations of an efficient serial algorithm on the Intel Paragon. One SIMD algorithm is quite simple, but its cost is linear in the cache size. The two other SIMD algorithm are more complex, but have costs that are independent of the cache size. Both the second and third SIMD algorithms compute all stack distances; the second SIMD algorithm is completely general, whereas the third SIMD algorithm presumes and takes advantage of bounds on the range of reference tags. Both MIMD algorithm implemented on the Paragon are general, and compute all stack distances; they differ in one step that may affect their respective scalability. We assess the strengths and weaknesses of these algorithms as a function of problem size and characteristics, and compare their performance on traces derived from execution of three SPEC benchmark programs.

Mao, Weizhen, and David Nicol: *On k-ary n-cubes: theory and applications.* ICASE Report No. 94-88, November 16, 1994, 29 pages. Submitted to SIAM Journal of Computing.

Many parallel processing networks can be viewed as graphs called k-ary n-cubes, whose special cases include rings, hypercubes and toruses. In this paper, combinatorial properties of k-ary n-cubes are explored. In particular, the problem of characterizing the subgraph of a given number of nodes with the maximum edge count is studied. These theoretical results are then used to compute a lower bounding function in branch-and-bound partitioning algorithms and to establish the optimality of some irregular partitions.

Van Rosendale, John: *Floating shock fitting via lagrangian adaptive meshes.* ICASE Report No.94-89 December 19, 1994, 24 pages. Submitted to 1995 AIAA Harbor Island CFD Conference.

In recent work we have formulated a new approach to compressible flow simulation, combining the advantages of shock-fitting and shock-capturing. Using a cell-centered Roe scheme discretization on unstructured meshes, we warp the mesh while marching to steady state, so that mesh edges align with shocks and other discontinuities. This new algorithm, the Shock-fitting Lagrangian Adaptive Method (SLAM) is, in effect, a reliable shock-capturing algorithm which yields shock-fitted accuracy at convergence.

Shock-capturing algorithms like this, which warp the mesh to yield shock-fitted accuracy, are new and relatively untried. However, their potential is clear. In the context of sonic booms, accurate calculation of near-field sonic boom signatures is critical to the design of the High Speed Civil Transport (HSCT). SLAM should allow computation of accurate N-wave pressure signatures on comparatively coarse meshes, significantly enhancing our ability to design low-boom configurations for high speed aircraft.

Tanveer, Saleh, and Giovani Vasconcelos: *Time-evolving bubbles in two-dimensional Stokes flow.* ICASE Report No. 94-90, December 7, 1994, 30 pages. To be submitted to the Journal of Fluid Mechanics.

A general class of exact solutions is presented for a time evolving bubble in a two-dimensional slow viscous flow in the presence of surface tension. These solutions can describe a bubble in a linear shear flow as well as an expanding or contracting bubble in an otherwise quiescent flow. In the case of expanding bubbles, the solutions have a simple behavior in the sense that for essentially arbitrary initial shapes the bubble will asymptote an expanding circle. Contracting bubbles, on the other hand, can develop narrow structures ('near-cusps') on the interface and may undergo 'break up' before all the bubble-fluid is completely removed. The mathematical structure underlying the existence of these exact solutions is also investigated.

Costiner, Sorin, and Shlomo Ta'asan: *Multigrid techniques for nonlinear eigenvalue problems; solutions of a nonlinear Schrödinger eigenvalue problem in 2D and 3D.* ICASE Report No. 94-91, November 16, 1994, 40 pages. To be submitted to Physical Review E, August 1994.

Algorithms for nonlinear eigenvalue problems (EP), often require solving selfconsistently a large number of EP. Convergence difficulties may occur if the solution is not sought in a right neighborhood; if global constraints have to be satisfied; and if close or equal eigenvalues are present. Multigrid (MG) algorithms for nonlinear problems and for EP obtained from discretizations of partial differential EP, have often shown to be more efficient than single level algorithms.

This paper presents MG techniques for nonlinear EP and emphasizes an MG algorithm for a nonlinear Schrödinger EP. The algorithm overcomes the mentioned difficulties combining the following techniques: an MG projection coupled with backrotations for separation of solutions and treatment of difficulties related to clusters of close and equal eigenvalues; MG subspace continuation techniques for the treatment of the nonlinearity; an MG simultaneous treatment of the eigenvectors at the same time with the nonlinearity and with the global constraints. The simultaneous MG techniques reduce the large number of selfconsistent iterations to only a few or one MG simultaneous iteration and keep the solutions in a right neighborhood where the algorithm converges fast.

Computational examples for the nonlinear Schrödinger EP in 2D and 3D, presenting special computational difficulties, which are due to the nonlinearity and to the equal and closely clustered eigenvalues, are demonstrated. For these cases, the algorithm requires $O(qN)$ operations for the calculation of q eigenvectors of size N and for the corresponding eigenvalues. One MG simultaneous

cycle per fine level was performed. The total computational cost is equivalent to only a few Gauss-Seidel relaxations per eigenvector. An asymptotic convergence rate of 0.15 per MG cycle is attained.

Zhou, Ye, Alexander Praskovsky, and Steven Oncley: *On the lighthill relationship and sound generation from isotropic turbulence.* ICASE Report No. 94-92, November 16, 1994, 18 pages. To appear in Theoretical and Computational Fluid Dynamics.

In 1952, Lighthill (*Proc. Roy. Soc.* **A211**, (1952)) developed a theory for determining the sound generated by a turbulent motion of a fluid. With some statistical assumptions, Proudman (*Proc. Roy. Soc.* **A211**, (1952)) applied this theory to estimate the acoustic power of isotropic turbulence. Recently, Lighthill established a simple relationship that relates the fourth-order retarded time and space covariance of his stress tensor to the corresponding second-order covariance and the turbulent flatness factor, without making statistical assumptions for a homogeneous turbulence. Lilley (Submitted to *Theoretical and Computational Fluid Dynamics*, 1993) revisited the Proudman's work and applied the Lighthill relationship to directly evaluate the radiated acoustic power from isotropic turbulence. After choosing the time separation dependence in the two-point velocity time and space covariance based on the insights gained from direct numerical simulations, Lilley concluded that the Proudman constant is determined by the turbulent flatness factor and the second-order spatial velocity covariance. In order to estimate the Proudman constant at high Reynolds numbers, we analyzed a unique data set of measurements in a large wind tunnel and atmospheric surface layer that covers a range of the Taylor microscale based Reynolds number $2.0 \times 10^3 \leq R_\lambda \leq 12.7 \times 10^3$. Our measurements demonstrate that the Lighthill relationship is a good approximation, providing additional support to Lilley's approach. The flatness factor is found between $2.7 - 3.3$ and the second order spatial velocity covariance is obtained. Based on these experimental data, the Proudman constant is estimated to be $0.68 - 3.68$.

Tanveer, Saleh: *A Note on singularities of the 3-D Euler equation.* ICASE Report No. 94-93, December 7, 1994, 13 pages. Submitted to Physics of Fluids.

In this paper, we consider analytic initial conditions with finite energy, whose complex spatial continuation is a superposition of a smooth background flow and a singular field. Through explicit calculation in the complex plane, we show that under some assumptions, the solution to the 3-D Euler equation ceases to be analytic in the real domain in finite time.

Zurigat, Y. H., and M. R. Malik: *Effect of crossflow on Görtler instability in incompressible boundary layers.* ICASE Report No. 94-94, December 7, 1994, 19 pages. To be submitted to Physics of Fluids.

Linear stability theory is used to study the effect of crossflow on Görtler instability in incompressible boundary layers. The results cover a wide range of sweep angle, pressure gradient, and wall curvature parameters. It is shown that the crossflow stabilizes Görtler disturbances by reducing the maximum growth rate and shrinking the unstable band of spanwise wave numbers. On

the other hand, the effect of concave wall curvature on crossflow instability is destabilizing. Calculations show that the changeover from Görtler to crossflow instabilities is a function of Görtler number, pressure gradient and sweep angle. The results demonstrate that Görtler instability may still be relevant in the transition process on swept wings even at large angles of sweep if the pressure gradient is sufficiently small. The influence of pressure gradient and sweep can be combined by defining a crossflow Reynolds number. Thus, the changeover from Görtler to crossflow instability takes place at some critical crossflow Reynolds number whose value increases with Görtler number.

Nicol, David M., and Philip Heidelberger: *On extending parallelism to serial simulators.* ICASE Report No. 94-95, December 15, 1994, 17 pages. Submitted to Workshop on Parallel and Distributed Simulation.

This paper describes an approach to discrete event simulation modeling that appears to be effective for developing portable and efficient parallel execution of models of large distributed systems and communication networks. In this approach, the modeler develops sub-models using an existing sequential simulation modeling tool, using the full expressive power of the tool. A set of modeling language extensions permit automatically synchronized communication between sub-models; however, the automation requires that any such communication must take a non-zero amount of simulation time. Within this modeling paradigm, a variety of conservative synchronization protocols can transparently support conservative execution of sub-models on potentially different processors. A specific implementation of this approach, UPS (Utilitarian Parallel Simulator), is described, along with performance results on the Intel Paragon.

Balling, R. J., and J. Sobieszczanski-Sobieski: *An Algorithm for solving the system-level problem in multilevel optimization.* ICASE Report No. 94-96, December 15, 1994, 25 pages. Proceedings of AIAA Paper No. 94-4333 and proceedings of AIAA/USAF/NASA/ISSMO Symposium on Multidisciplinary Analysis and Optimization.

A multilevel optimization approach which is applicable to nonhierarchic coupled systems is presented. The approach includes a general treatment of design (or behavior) constraints and coupling constraints at the discipline level through the use of norms. Three different types of norms are examined–the max norm, the Kreisselmeier-Steinhauser (KS) norm, and the l_p norm. The max norm is recommended. The approach is demonstrated on a class of hub frame structures which simulate multidisciplinary systems. The max norm is shown to produce system-level constraint functions which are non-smooth. A cutting-plane algorithm is presented which adequately deals with the resulting corners in the constraint functions. The algorithm is tested on hub frames with increasing number of members (which simulate disciplines), and the results are summarized.

Joslin, Ronald D., Gordon Erlebacher, and M. Yousuff Hussaini: *Active control of instabilities in laminar boundary-layer flow - part I: an overview.* ICASE Report No. 94-97, December 16, 1994, 15 pages. Submitted to Journal of Fluids Engineering.

This paper (the first in a series) focuses on using active-control methods to maintain laminar flow in a region of the flow in which the natural instabilities, if left unattended, lead to turbulent flow. The authors review previous studies that examine wave cancellation (currently the most prominent method) and solve the unsteady, nonlinear Navier-Stokes equations to evaluate this method of controlling instabilities. It is definitely shown that instabilities are controlled by the linear summation of waves (i.e., wave cancellation). Although a mathematically complete method for controlling arbitrary instabilities has of previous studies are important steps for providing an independent verification of those studies, for establishing a framework for subsequent work which will involve automated transition control, and for detailing the phenomena by-which the automated studies can be used to expand knowledge of flow control.

Joslin, Ronald D., R. A. Nicolaides, Gordon Erlebacher, M. Yousuff Hussaini, and Max D. Gunzburger: *Active control of instabilities in laminar boundary-layer flow - part II: use of sensors and spectral controller.* ICASE Report No. 94-98, December 16, 1994, 11 pages. Submitted to AIAA Journal.

This study focuses on the suppression of instability growth using an automated active-control technique. The evolution of 2D disturbances that are spatially growing in a flat-plate boundary layer are computed with a spatial DNS code. A controller receives wall sensor information (pressure or shear) as input and provides a signal that controls an actuator response as output. The control law assumes that wave cancellation is valid. The results indicate that a measure of wave cancellation can be obtained for small- and large-amplitude instabilities without feedback; however, feedback is required to optimize the control amplitude and phase for exact wave cancellation.

Bryan, Kurt, and Lester F. Caudill, Jr.: *An Inverse problem in thermal imaging.* ICASE Report No. 94-99, December 15, 1994, 34 pages. Submitted to SIAM Applied Mathematics.

This paper examines uniqueness and stability results for an inverse problem in thermal imaging. The goal is to identify an unknown boundary of an object by applying a heat flux and measuring the induced temperature on the boundary of the sample. The problem is studied both in the case in which one has data at every point on the boundary of the region and the case in which only finitely many measurements are available. An inversion procedure is developed and used to study the stability of the inverse problem for various experimental configurations.

Balling, R. J., and J. Sobieszczanski-Sobieski: *Optimization of coupled systems: a critical overview of approaches.* ICASE Report No. 94-100, December 15, 1994, 32 pages. AIAA Paper No. 94-4330 and Proceedings AIAA/USAF/NASA/ISSMO Symposium on Multidisciplinary Analysis and Optimization, Panama City, FL, September 7-9, 1994.

A unified overview is given of problem formulation approaches for the optimization of multidisciplinary coupled systems. The overview includes six fundamental approaches upon which a large number of variations may be made. Consistent approach names and a compact approach notation are given. The approaches are formulated to apply to general nonhierarchic systems. The approaches are compared both from a computational viewpoint and a managerial viewpoint. Opportunities for parallelism of both computation and manpower resources are discussed. Recommendations regarding the need for future research are advanced.

Bertolotti, Fabio P., and Ronald D. Joslin: *Effect of far-field boundary conditions on boundary-layer transition.* ICASE Report No. 94-101, December 19, 1994, 18 pages. Submitted to the Journal of Computational Physics.

The effect of far-field boundary conditions on the evolution of a finite-amplitude two-dimensional wave in the Blasius boundary layer is assessed. With the use of the parabolized stability equations (PSE) theory for the numerical computations, either asymptotic, Dirichlet, Neumann or mixed boundary conditions are imposed at various distances from the wall. The results indicate that asymptotic and mixed boundary conditions yield the most accurate mean-flow distortion and unsteady instability modes in comparison with the results obtained with either Dirichlet or Neumann conditions.

Hu, F. Q., M. Y. Hussaini, and J. Manthey: *Low-dissipation and -dispersion Runge-Kutta schemes for computational acoustics.* ICASE Report No. 94-102, December 19, 1994, 25 pages. Submitted to Journal of Computational Physics.

In this paper, we investigate accurate and efficient time advancing methods for computational acoustics, where non-dissipative and non-dispersive properties are of critical importance. Our analysis pertains to the application of Runge-Kutta methods to high-order finite difference discretization. In many CFD applications, multi-stage Runge-Kutta schemes have often been favored for their low storage requirements and relatively large stability limits. For computing acoustic waves, however, the stability consideration alone is not sufficient, since the Runge-Kutta schemes entail both dissipation and dispersion errors. The time step is now limited by the tolerable dissipation and dispersion errors in the computation. In the present paper, it is shown that if the traditional Runge-Kutta schemes are used for time advancing in acoustic problems, time steps greatly smaller than that allowed by the stability limit are necessary. Low-Dissipation and -Dispersion Runge-Kutta (LDDRK) schemes are proposed, based on an optimization that minimizes the dissipation and dispersion errors for wave propagation. order Optimizations of both single-step and two-step

alternating schemes are considered. The proposed LDDRK schemes are remarkably more efficient than the classical Runge-Kutta schemes for acoustic computations. Moreover, low storage implementations of the optimized schemes are discussed. Special issues of implementing numerical boundary conditions in the LDDRK schemes are also addressed.

Papageorgiou, Demetrios T.: *On the breakup of viscous liquid threads.* ICASE Report No. 95-1, January 11, 1995, 49 pages. Submitted to Physics of Fluids.

A one-dimensional model evolution equation is used to describe the nonlinear dynamics that can lead to the breakup of a cylindrical thread of Newtonian fluid when capillary forces drive the motion. The model is derived from the Stokes equations by use of rational asymptotic expansions and under a slender jet approximation. The equations are solved numerically and the jet radius is found to vanish after a finite time yielding breakup. The slender jet approximation is valid throughout the evolution leading to pinching. The model admits self-similar pinching solutions which yield symmetric shapes at breakup. These solutions are shown to be the ones selected by the initial boundary value problem, for general initial conditions. Further more, the terminal state of the model equation is shown to be identical to that predicted by a theory which looks for singular pinching solutions directly from the Stokes equations without invoking the slender jet approximation throughout the evolution. It is shown quantitatively, therefore, that the one-dimensional model gives a consistent terminal state with the jet shape being locally symmetric at breakup. The asymptotic expansion scheme is also extended to include unsteady and inertial forces in the momentum equations to derive an evolution system modelling the breakup of Navier-Stokes jets. The model is employed in extensive simulations to compute breakup times for different initial conditions; satellite drop format ion is also supported by the model and the dependence of satellite drop volumes on initial conditions is studied.

Brunstrom, Anna, Scott T. Leutenegger, and Rahul Simha: *Experimental evaluation of dynamic data allocation strategies in a distributed database with changing workloads.* ICASE Report No. 95-2, January 11, 1995, 17 pages. Submitted to SIGMOD '95.

Traditionally, allocation of data in distributed database management systems has been determined by off-line analysis and optimization. This technique works well for static database access patterns, but is often inadequate for frequently changing workloads. In this paper we address how to dynamically reallocate data for partionable distributed databases with changing access patterns. Rather than complicated and expensive optimization algorithms, a simple heuristic is presented and shown, via an implementation study, to improve system throughput by 30% in a local area network based system. Based on artificial wide area network delays, we show that dynamic reallocation can improve system throughput by a factor of two and a half for wide area networks. We also show that individual site load must be taken into consideration when reallocating data, and provide a simple policy that incorporates load in the reallocation decision.

Venkatakrishnan, V.: *A perspective on unstructured grid flow solvers.* ICASE Report No. 95-3, February 15, 1995, 37 pages. Invited paper presented at the 33rd AIAA Aerospace Sciences Conference; Submitted to AIAA Journal.

This survey paper assesses the status of compressible Euler and Navier-Stokes solvers on unstructured grids. Different spatial and temporal discretization options for steady and unsteady flows are discussed. The integration of these components into an overall framework to solve practical problems is addressed. Issues such as grid adaptation, higher order methods, hybrid discretizations and parallel computing are briefly discussed. Finally, some outstanding issues and future research directions are presented.

Nicolaides, R. A., and X. Wu: *Covolume solutions of three dimensional div-curl equations.* ICASE Report No. 95-4, February 2, 1995, 13 pages. Submitted to the SIAM Journal of Numerical Analysis.

Delaunay-Voronoi mesh systems provide a generalization of the classical rectangular staggered meshes to unstructured meshes. In this work, it is shown how such "covolume" discretizations may be applied to div-curl systems in three dimensions. Error estimates are proved and confirmed by a numerical illustration.

Verhaagen, N. G., L. N. Jenkins, S. B. Kern, and A. E. Washburn: *A study of the vortex flow over 76/40-deg double-delta wing.* ICASE Report No. 95-5, February 15, 1995, 33 pages. AIAA Paper 95-0560; Presented at 33rd AIAA Aerospace Sciences Meeting.

A low-speed wind-tunnel study of the flow about a 76/40-deg double-delta wing is described for angles of attack ranging from -10 to 25 deg and Reynolds numbers ranging from 0.5 to 1.5 Million. The study was conducted to provide data for the purpose of understanding the vortical flow behavior and for validating Computational Fluid Dynamics methods. Flow visualization tests have provided insight into the effect of the angle of attack and Reynolds number on the vortex-dominated flow both on and off of the surface of the double-delta wing. Upper surface pressure recordings from pressure orifices and Pressure Sensitive Paint have provided data on the pressures induced by the vortices. Flowfield surveys were carried out at an angle of attack of 10 deg by using a thin 5-hole probe. Numerical solutions of the compressible thin-layer Navier-Stokes equations were conducted and compared to the experimental data.

Criminale, W. O., T. L. Jackson, and D. G. Lasseigne: *The initial-value problem for viscous channel flows.* ICASE Report No. 95-6, February 15, 1995, 29 pages. Submitted to the Journal of Fluid Mechanics.

Plane viscous channel flows are perturbed and the ensuing initial-value problems are investigated in detail. Unlike traditional methods where traveling wave normal modes are assumed

for solution, this work offers a means whereby completely arbitrary initial input can be specified without having to resort to eigenfunction expansions. The full temporal behavior, including both early time transients and the long time asymptotics, can be determined for any initial disturbance. Effects of three-dimensionality can be assessed. The bases for the analysis are: (a) linearization of the governing equations; (b) Fourier decomposition in the spanwise and streamwise directions of the flow and; (c) direct numerical integration of the resulting partial differential equations. All of the stability data that are known for such flows can be reproduced. Also, the optimal initial conditions can be determined in a straight forward manner and such optimal conditions clearly reflect transient growth data that is easily determined by a rational choice of a basis for the initial conditions. Although there can be significant transient growth for subcritical values of the Reynolds number using this approach, it does not appear possible that arbitrary initial conditions will lead to the exceptionally large transient amplitudes that have been determined by optimization of normal modes. The approach is general and can be applied to other classes of problems where only a finite discrete spectrum exists, such as the boundary layer for example.

Baker, Gregory, Michael Siegel, and Saleh Tanveer: *A well-posed numerical method to track isolated conformal map singularities in hele-shaw flow.* ICASE Report No. 95-7, February 17, 1995, 50 pages. Submitted to the Journal of Computational Physics.

We present a new numerical method for calculating an evolving 2-D Hele-Shaw interface when surface tension effects are neglected. In the case where the flow is directed from the less viscous fluid into the more viscous fluid, the motion of the interface is ill-posed; small deviations in the initial condition will produce significant changes in the ensuing motion. This situation is disastrous for numerical computation, as small round-off errors can quickly lead to large inaccuracies in the computed solution. Our method of computation is most easily formulated using a conformal map from the fluid domain into a unit disk. The method relies on analytically continuing the initial data and equations of motion into the region exterior to the disk, where the evolution problem becomes well-posed. The equations are then numerically solved in the extended domain. The presence of singularities in the conformal map outside of the disk introduces specific structures along the fluid interface. Our method can explicitly track the location of isolated pole and branch point singularities, allowing us to draw connections between the development of interfacial patterns and the motion of singularities as they approach the unit disk. In particular, we are able to relate physical features such as finger shape, side-branch formation, and competition between fingers to the nature and location of the singularities. The usefulness of this method in studying the formation of topological singularities (self-intersections of the interface) is also pointed out.

Costiner, Sorin, and Shlomo Ta'asan: *Separation analysis, a tool for analyzing multigrid algorithms.* ICASE Report No. 95-8, February 17, 1995, 20 pages. To be submitted to the Journal of Computational Physics.

The separation of vectors by multigrid (MG) algorithms is applied to the study of convergence and to the prediction of the performance of MG algorithms. The separation operator for a two

level cycle algorithm is derived. It is used to analyze the efficiency of the cycle when mixing of eigenvectors occurs. In particular cases the separation analysis reduces to Fourier type analysis. The separation operator of a two level cycle for a Schrödinger eigenvalue problem, is derived and analyzed in a Fourier basis. Separation analysis gives information on how to chose relaxations and inter-level transfers. Separation analysis is a tool for analyzing and designing algorithms, and for optimizing their performance.

Ciardo, Gianfranco: *Discrete-time markovian stochastic petri nets.* ICASE Report No. 95-9, March 2, 1995, 20 pages. To appear in Numerical Solution of Markov Chains '95.

We revisit and extend the original definition of discrete-time stochastic Petri nets, by allowing the firing times to have a "defective discrete phase distribution". We show that this formalism still corresponds to an underlying discrete-time Markov chain. The structure of the state for this process describes both the marking of the Petri net and the phase of the firing time for of each transition, resulting in a large state space. We then modify the well-known power method to perform a transient analysis even when the state space is infinite, subject to the condition that only a finite number of states can be reached in a finite amount of time. Since the memory requirements might still be excessive, we suggest a bounding technique based on truncation.

Sidilkover, D., and P. L. Roe: *Unification of some advection schemes in two dimensions.* ICASE Report No. 95-10, March 2, 1995, 28 pages. Submitted to Math. Comp.

In this paper a relationship between two approaches towards construction of genuinely two-dimensional upwind advection schemes is established. One of these approaches is of the control volume type applicable on structured cartesian meshes. It resulted (see [14], [15]) in the compact high resolution schemes capable of maintaining second order accuracy in both homogeneous and inhomogeneous cases. Another one is the fluctuation splitting approach (see [11], [3], [12], [17]), which is well suited for triangular (and possibly) unstructured meshes. Understanding the relationship between these two approaches allows us to formulate here a new fluctuation splitting high resolution (i.e. possible use of artificial compression, while maintaining positivity property) scheme. This scheme is shown to be linearity preserving in inhomogeneous as well as homogeneous cases.

Wesseling, Pieter: *Introduction to multigrid methods.* ICASE Report No. 95-11, March 2, 1995, 136 pages.

These notes were written for an introductory course on the application of multigrid methods to elliptic and hyperbolic partial differential equations for engineers, physicists and applied mathematicians. The use of more advanced mathematical tools, such as functional analysis, is avoided. The course is intended to be accessible to a wide audience of users of computational methods. We restrict ourselves to finite volume and finite difference discretization. The basic principles are

given. Smoothing methods and Fourier smoothing analysis are reviewed. The fundamental multigrid algorithm is studied. The smoothing and coarse grid approximation properties are discussed. Multigrid schedules and structured programming of multigrid algorithms are treated. Robustness and efficiency are considered.

Salas, Manuel D., and Angelo Iollo: *Entropy jump across an inviscid shock wave.* ICASE Report No. 95-12, March 2, 1995, 16 pages. To be submitted to Theoretical and Computational Fluid Dynamics.

The shock jump conditions for the Euler equations in their primitive form are derived by using generalized functions. The shock profiles for specific volume, speed, and pressure are shown to be the same, however density has a different shock profile. Careful study of the equations that govern the entropy shows that the inviscid entropy profile has a local maximum within the shock layer. We demonstrate that because of this phenomenon, the entropy propagation equation cannot be used as a conservation law.

Girimaji, Sharath S., and Ye Zhou: *Spectrum and energy transfer in steady Burgers turbulence.* ICASE Report No. 95-13, March 17, 1995, 19 pages. Submitted to Physics Letters A.

The spectrum, energy transfer and spectral interactions in steady Burgers turbulence are studied using numerically generated data. The velocity field is initially random and the turbulence is maintained steady by forcing the amplitude of a band of low wavenumbers to be invariant in time, while permitting the phase to change as dictated by the equation. The spectrum, as expected, is very different from that of Navier-Stokes turbulence. It is demonstrated, for the first time, that the far range of the spectrum scales as predicted by Burgers (1950, 1974). Despite the difference in their spectra, in matters of the spectral energy transfer and triadic interactions Burgers turbulence is similar to Navier-Stokes turbulence.

Mavriplis, D. J.: *Multigrid solution strategies for adaptive meshing problems.* ICASE Report No. 95-14, March 17, 1995, 39 pages. To appear in Proceedings of Adaptive Mesh Workshop, as NASA Conference Publication.

This paper discusses the issues which arise when combining multigrid strategies with adaptive meshing techniques for solving steady-state problems on unstructured meshes. A basic strategy is described, and demonstrated by solving several inviscid and viscous flow cases. Potential inefficiencies in this basic strategy are exposed, and various alternate approaches are discussed, some of which are demonstrated with an example. Although each particular approach exhibits certain advantages, all methods have particular drawbacks, and the formulation of a completely optimal strategy is considered to be an open problem.

Speziale, Charles G.: *A review of Reynolds stress models for turbulent shear flows.* ICASE Report No. 95-15, March 17, 1995, 44 pages. To appear in the Proceedings of the Twentieth Symposium on Naval Hydrodynamics.

A detailed review of recent developments in Reynolds stress modeling for incompressible turbulent shear flows is provided. The mathematical foundations of both two-equation models and full second-order closures are explored in depth. It is shown how these models can be systematically derived for two-dimensional mean turbulent flows that are close to equilibrium. A variety of examples are provided to demonstrate how well properly calibrated versions of these models perform for such flows. However, substantial problems remain for the description of more complex turbulent flows where there are large departures from equilibrium. Recent efforts to extend Reynolds stress models to non-equilibrium turbulent flows are discussed briefly along with the major modeling issues relevant to practical Naval Hydrodynamics applications.

Morano, E., and D. J. Mavriplis: *Implementation of a parallel unstructured Euler solver on the CM-5.* ICASE Report No. 95-16, March 17, 1995, 22 pages. To be submitted to the International Journal of Computational Fluid Dynamics.

An efficient unstructured 3D Euler solver is parallelized on a Thinking Machine Corporation Connection Machine 5, distributed memory computer with vectorizing capability. In this paper, the SIMD strategy is employed through the use of the CM Fortran language and the CMSSL scientific library. The performance of the CMSSL mesh partitioner is evaluated and the overall efficiency of the parallel flow solver is discussed.

Geer, James, and Nana Saheb Banerjee: *Exponentially accurate approximations to piece-wise smooth periodic functions.* ICASE Report No. 95-17, March 21, 1995, 40 pages. To be submitted to SIAM Journal on Scientific Computation.

A family of simple, periodic basis functions with "built-in" discontinuities are introduced, and their properties are analyzed and discussed. Some of their potential usefulness is illustrated in conjunction with the Fourier series representation of functions with discontinuities. In particular, it is demonstrated how they can be used to construct a sequence of approximations which converges *exponentially in the maximum norm* to a piece-wise smooth function. The theory is illustrated with several examples and the results are discussed in the context of other sequences of functions which can be used to approximate discontinuous functions.

Kopriva, David A., and John H. Kolias: *A conservative staggered-grid Chebyshev multidomain method for compressible flows.* ICASE Report No. 95-18, March 23, 1995, 48 pages. Submitted to the Journal of Computational Physics.

We present a new multidomain spectral collocation method that uses staggered grids for the solution of compressible flow problems. The solution unknowns are defined at the nodes of a Gauss quadrature rule. The fluxes are evaluated at the nodes of a Gauss-Lobatto rule. The method is conservative, free-stream preserving and exponentially accurate. A significant advantage of the method is that subdomain corners are not included in the approximation, making solutions in complex geometries easier to compute.

Zhou, Ye: *Classical closure theory and Lam's interpretation of ϵ-RNG.* ICASE Report No. 95-19, March 28, 1995, 10 pages. To appear in Physical Review E.

Lam's phenomenological ϵ-renormalization group (RNG) model is quite different from the other members of that group. It does not make use of the correspondence principle and the ϵ-expansion procedure. In this report, we demonstrate that Lam's ϵ-RNG model [*Phys. Fluids A*, **4**, 1007 (1992)] is essentially the physical space version of the classical closure theory [Leslie and Quarini, *J. Fluid Mech.*, **91**, 65 (1979)] in spectral space and consider the corresponding treatment of the eddy viscosity and energy backscatter.

Iollo, Angelo, and Manuel D. Salas: *Contribution to the optimal shape design of two-dimensional internal flows with embedded shocks.* ICASE Report No. 95-20, March 28, 1995, 22 pages. To be submitted to the Journal of Computational Physics.

We explore the praticability of optimal shape design for flows modeled by the Euler equations. We define a functional whose minimum represents the optimality condition. The gradient of the functional with respect to the geometry is calculated with the Lagrange multipliers, which are determined by solving a *costate* equation. The optimization problem is then examined by comparing the performance of several gradient-based optimization algorithms. In this formulation, the flow field can be computed to an arbitrary order of accuracy. Finally, some results for internal flows with embedded shocks are presented, including a case for which the solution to the inverse problem does not belong to the design space.

ICASE COLLOQUIA

October 1, 1994 - March 31, 1995

Name/Affiliation/Title	Date
Farhat, Charbel, University of Colorado at Boulder "Mixed Explicit/Implicit Time Integration of Coupled Aeroelastic Problems: Three-Field Formulation, Geometric Conservation and Distributed Solution"	October 6
Babb, Robert, University of Denver "Semi-Automatic Re-Engineering of Scientific Programs for Large-Grain Parallelism"	October 21
Uselton, Samuel, Computer Sciences Corp., NASA Ames Research Center "Visualization of Large Computational Aeroscience Data Sets"	October 24
Hussain, Fazle, University of Houston "New Perspectives on Vortex Dynamics, Coherent Structures, and Related Turbulence Physics"	November 3
Kelley, C. Timothy, North Carolina State University "Implicit Filtering and Noisy Optimization Problems"	November 15
Chen, Ching-Jen, Florida State University "Finite Analytic Numerical Simulation and Quantitative Flow Visualization of Complex Flows"	November 17
Wakayama, Sean, Stanford University "Lifting Surface Design using Multidisciplinary Optimization"	Novembr 18
McInnes, Lois Curfman, Argonne National Laboratory "Portable, Extensible, and Scalable Optimization Libraries"	November 21
Xue, Zachary, North Carolina State University "Mesh-Independence of GMRES for Integral Equations"	November 28
Abarbanel, Saul, Tel-Aviv University, Israel "Error Bounds on Long Time Integration of Parabolic P.D.E.s"	November 29
Kearsley, Anthony José, Rice University "Using Parallelism to Solve a Class of Shape Optimization Problems"	December 5

Name/Affiliation/Title	Date
Purtell, Patrick, Office of Naval Research "Turbulence Research Program at the Office of Naval Research"	December 14
Haj-Harari, Hossein, University of Virginia "Effect of Viscosity on the Wall-Pressure Spectrum of a Flat-Plate Turbulent Boundary Layer"	December 16
Özturan, Can, Rensselaer Polytechnic Institute "Mesh Migration and Dynamic Load Balancing Algorithms for 3D Unstructured Adaptive Finite Element Meshes"	December 19
Havener, George, Micro Craft Technology-AEDC Operations "Computational Flow Imaging (CFI)"	December 21
Lewis, Robert Michael, Positron Corporation, Houston, TX "Nonlinear Programming Approaches to the Optimization of Complex, Coupled Systems"	January 9
Torczon, Virginia, Rice University "Pattern Search Methods for Nonlinear Optimization"	January 10
Verhaagen, Nicholas, Delft University of Technology, The Netherlands "The Vortical Flow over Double-Delta Wings"	January 18
Thakur, Rajeev, Syracuse University "Runtime Support for Out-of-Core Computations"	January 20
Kennedy, Ken, Rice University "Software Support for High Performance Fortran User"	January 23
Chakravarthy, Sukumar, University of California, Los Angeles "Interior Boundary Conditions for Inviscid Flows"	January 24
Chakravarthy, Sukumar, University of California, Los Angeles "'Commercialization' of Computational Fluid Dynamics: Examples with Lessons Learned"	January 25
Dwoyer, Douglas, Research and Technology Group (RTG), NASA "Third Generation R & D in RTG"	January 27
Venkatakrishnan, V., ICASE "A Perspective on Unstructured Grid Flow Solvers"	January 27

Name/Affiliation/Title	Date
Ristorcelli, J. Ray, ICASE "A Pseudo-Sound Constitutive Relationship for the Dilatational Covariances in Compressible Turbulence: An Analytical Theory"	February 17
Abarbanel, Saul, Tel-Aviv, Israel "Bounded-Error Algorithms for Scientific Computations"	February 21
Cheung, Samson, NASA Ames Research Center "Aerodynamic Optimization and Challenges of Multidisciplinary Design on Parallel Machines"	February 23
Fridrich, Jiri, State University of New York "Removing Observational Uncertainty from Chaotic Dynamical Systems"	February 24
Krothapalli, Anjaneyulu, FAMU/FSU College of Engineering "On Supersonic Jet Mixing"	February 24
Wong, Sze-Ping, Elastic Reality Incorporated, Madison, WI "The Distribution of Singular Values in Preconditioning"	February 27
Polyak, Roman, George Mason University "Modified Barrier Functions and the 'Hot' Start Phenomenon in Constrained Optimization"	February 28
Cordner, David, and Nelson, Michael, ISD - Advanced Computer Systems Branch, NASA "Workstation Clustering in the HPCCP Program"	March 2
Agarwal, Ramesh K., Wichita State University "A Compact Higher-Order Finite-Volume Time-Domain/Frequency-Domain Method for Aeroacoustics"	March 3
Allan, Brian, University of California, Berkeley "Multidisciplinary Coupling of Fluids, Controls, and Dynamics"	March 9
Xu, Jinchao, The Pennsylvania State University "Optimal Multigrid Methods for General Unstructured Grids"	March 10
Golubchik, Leana, University of California, Los Angeles "On Efficient Use of Resources in On-Demand Multimedia Storage Servers"	March 15

Name/Affiliation/Title	Date
Young, David P., Boeing Computer Services and the Boeing Commercial Airplane Company "Some Technical Points on Design and Optimization"	March 16
Olds, John R., North Carolina State University "MDO Applications and Research in Space Launch Vehicle Conceptual Design"	March 17
Geer, James, Binghamton University "Exponentially Convergent Approximations using Simple Periodic Basis Functions with Discontinuities"	March 21
Hayder, M. Ehtesham, ICOMP, NASA Lewis Research Center "A Scalability and Communication Study of Jet Noise Computations on Parallel Computers"	March 24
Bonnet, Jean Paul, CNRS and University of Poitiers, France "Some Challenging Problems in Predicting Experimental Results in Supersonic Turbulent Free Shear Flows"	March 27
Levin, Peter L., Worcester Polytechnic Institute "Multiscale Solutions of BEM Equations for Electroquasistatic Systems"	March 27
Roh, Lucas, Colorado State University "Multithreading: Efficiently Exploiting Parallelism"	March 31
Soubbaramayer, CEA-SACLAY, France "Fluid Mechanics in R & D on Uranium Enrichment"	March 31

OTHER ACTIVITIES

On October 3-4, 1994, ICASE and NASA LaRC co-sponsored an Industry Roundtable at the Williamsburg Hospitality House in Williamsburg, VA. The objective of this Roundtable was to expose ICASE/LaRC scientists to industrial research agendas and to acquaint industry with the capabilities and technology available at ICASE/LaRC. There were 200 attendees, and an ICASE interim report was published.

On October 24-26, 1994, ICASE and NASA LaRC co-sponsored a Workshop on Benchmark Problems in Computational Aeroacoustics at the Holiday Inn in Hampton, VA. This Workshop was a sequel to the first Workshop on Computational Aeroacoustics held in April 1992. The objective of this Workshop was to evaluate the accuracy and efficiency of numerical methods in a systematic fashion as applied to certain elementary testbed problems which characterize the difficulties of computational aeroacoustics. There were 86 attendees, and a NASA Conference Proceedings was published.

A Workshop on Adaptive Grid Methods was held November 7-9, 1994 at the Radisson Hotel in Hampton, VA. The objective of this Workshop was to bring together experts in adaptive refinement to evaluate the state-of-the-art in adaptive grid methods, and to discuss current research. There were 59 attendees, and a NASA Conference Proceedings will be published.

On November 14, 1994, Dr. Richard T. Whitcomb, retired distinguished research associate at NASA LaRC and former head of Langley's Transonic Aerodynamics Branch, inaugurated the Eastman Jacobs Award with a lecture at the H.J.E. Reid Conference Center entitled "Research on Methods for Reducing the Aerodynamic Drag at Transonic Speeds." This award is sponsored biennial by ICASE and NASA Langley Research Center to recognize individuals who have distinguished themselves by making significant experimental or design contributions to aeronautical engineering and sciences.

A Workshop on Multidisciplinary Design Optimization co-sponsored by ICASE and NASA LaRC was held March 13-16, 1995 at the Holiday Inn in Hampton, VA. The objective of this Workshop was to acquaint researchers and practitioners in the field of complex coupled systems with the state-of-the-art in MDO, to identify practical and theoretical needs and opportunities, and to engage the participants in discussions about the issues in MDO. There were 110 attendees. The proceedings will be published by SIAM.

ICASE STAFF

I. ADMINISTRATIVE

M. Yousuff Hussaini, Director. Ph.D., Mechanical Engineering, University of California, 1970.

Linda T. Johnson, Office and Financial Administrator

Etta M. Blair, Accounting Supervisor

Barbara A. Cardasis, Administrative Secretary

Shannon L. Keeter, Technical Publications Secretary

Rachel A. Lomas, Payroll and Accounting Clerk

Shelly D. Millen, Executive Secretary/Visitor Coordinator

Emily N. Todd, Conference Manager

Gwendolyn W. Wesson, Contract Accounting Clerk

Leon M. Clancy, System Manager

Sharon S. Paulson, Assistant System Manager, through December 1994.

Avik Banerjee, System Operator

II. SCIENCE COUNCIL

Ivo Babuska, Professor, Institute for Physical Science & Technology, University of Maryland.

Geoffrey Fox, Director, Northeast Parallel Architectural Center, Syracuse University.

Dennis Gannon, Professor, Center for Innovative Computer Applications, Indianna University.

Ashwani Kapila, Professor, Department of Mathematics and Science, Rensselaer Polytechnic Institute.

James P. Kendall, Jet Propulsion Laboratory (Convener).

Heinz-Otto Kreiss, Professor, Department of Mathematics, University of California at Los Angeles.

Sanjoy Mitter, Professor of Electrical Engineering, Massachusetts Institute of Technology.

Steven A. Orszag, Professor, Program in Applied & Computational Mathematics, Princeton University.

Paul Rubbert, Unit Chief, Boeing Commercial Airplane Group.

Ahmed Sameh, Department Head of Computer Science, University of Minnesota.

M. Y. Hussaini, Director, Institute for Computer Applications in Science and Engineering, NASA Langley Research Center (Ex-Officio).

IV. RESEARCH FELLOWS

Gordon Erlebacher - Ph.D., Plasma Physics, Columbia University, 1983. Computational Fluid Dynamics. (November 1989 to November 1999)

Piyush Mehrotra - Ph.D., Computer Science, University of Virginia, 1982. Programming Languages for Multiprocessor Systems. (January 1991 to September 1999)

IV. SENIOR STAFF SCIENTISTS

Thomas Jackson - Ph.D., Mathematics, Rensselaer Polytechnic Institute, 1985. Fluid Mechanics. (January 1994 to January 1997)

Dimitri Mavriplis - Ph.D., Mechanical and Aerospace Engineering, Princeton University, 1988. Grid Techniques for Computational Fluid Dynamics. (February 1987 to August 1997)

Robert Rubinstein - Ph.D., Mathematics, Massachusetts Institute of Technology, 1972. Fluid Mechanics. (January 1995 to January 1998)

David Sidilkover - Ph.D., Applied Mathematics, The Weizmann Institute of Science, 1989. Applied & Numerical Mathematics. (November 1994 to October 1996)

V. Venkatakrishnan - Ph.D., Mechanical and Aerospace Engineering, Princeton University, 1987. Applied & Numerical Mathematics. (June 1993 to September 1997)

V. SCIENTIFIC STAFF

Eyal Arian - Ph.D., Applied Mathematics, The Weizmann Institute of Science, Israel, 1995. Applied & Numerical Mathematics. (October 1994 to September 1996)

David Banks - Ph.D., Computer Science, University of North Carolina, 1993. Computer Science [Visualization]. (April 1993 to August 1996)

Phillip M. Dickens - Ph.D., Computer Science, University of Virginia, 1992. Computer Science [System Software]. (January 1993 to August 1996)

Sharath S. Girimaji - Ph.D., Mechanical and Aerospace Engineering, Cornell University, 1990. Fluid Mechanics [Turbulence and Combustion]. (July 1993 to July 1995)

Matthew D. Haines - Ph.D., Computer Science, Colorado State University, 1993. Computer Science [System Software]. (August 1993 to August 1996)

Leland M. Jameson - Ph.D., Applied Mathematics, Brown University, 1993. Applied & Numerical Mathematics. (October 1993 to October 1995)

Jim E. Jones - Ph.D., Applied Mathematics, University of Colorado-Boulder, 1995. Computer Science [Parallel Multigrid Methods]. (March 1995 to March 1997)

Kwan-Liu Ma - Ph.D., Computer Science, University of Utah, 1993. Computer Science [Visualization]. (May 1993 to August 1996)

Roberto Marsilio - Ph.D., Aerospace Engineering, Politecnico di Torino, Italy, 1992. Applied & Numerical Mathematics [Three Dimensional Shock Fitting and Shock-Induced Noise]. (May 1994 to May 1996)

James Quirk - Ph.D., Computational Fluid Dynamics, Cranfield Institute of Technology, 1991. Applied & Numerical Mathematics [Adaptive Methods for Partial Differential Equations]. (June 1991 to December 1994)

J. Ray Ristorcelli - Ph.D., Mechanical and Aerospace Engineering, Cornell University, 1991. Fluid Mechanics [Turbulence Modeling]. (December 1992 to September 1996)

Moulay D. Tidriri - Ph.D., Fluid Mechanics, University of Paris IX, France, 1992. Computer Science [Domain Decomposition Methods and Parallel Solution of Partial Differential Equations]. (September 1994 to September 1995)

Linda F. Wilson - Ph.D., Electrical and Computer Engineering, University of Texas at Austin, 1994. Computer Science. [Investigating the Issues of Object Migration in Object-Oriented Performance Tools and Object-Oriented Modeling of Architectures for Joint Performance/Reliability Analysis]. (August 1994 - July 1996)

Ye Zhou - Ph.D., Physics, College of William and Mary, 1987. Fluid Mechanics [Turbulence Modeling]. (October 1992 to September 1996)

VI. SENIOR SYSTEMS ANALYST

Thomas W. Crockett - B.S., Mathematics, College of William & Mary, 1977. Computer Science [System Software for Parallel Computing, Computer Graphics, and Scientific Visualization]. (February 1987 to August 1997)

VI. VISITING SCIENTISTS

Francoise Bataille - Ph.D., Compressable Turbulence, Ecole Centrale de Lyon, France, 1994. Doctor, Laboratoire de Mecanique des Fluides et d'Acoustique, France. Fluid Mechanics [Direct Numerical Simulation of Compressible Turbulence]. (October 1994 to September 1995)

Barbara Chapman - M.S., Mathematics, University of Canterbury, Christchurch, New Zealand, 1985. Research Associate, Computer Science Department, University of Vienna. Computer Science. (February to March 1995)

Thierry Dubois - Ph.D., Numerical Analysis, Orsay University, Paris, France, 1993. Attache Temporaire d'Enseignement et de Recherche, Laboratoire d'Analyse Numerique, Universite de Paris-Sud, France. Applied & Numerical Mathematics. (October to December 1995)

Jacques Liandrat - Ph.D., Fluid Mechanics-Numerical Simulations, University of Toulouse, France, 1986. Research Scientist, I.M.S.T., Marseille, France. Applied & Numerical Mathematics [Numerical Algorithms for Nonlinear PDE Based on Wavelets and Transition Analysis Using the Wavelet Approach]. (October 1994)

Michael N. Macrossan - Ph.D., Rarefied Gas Dynamics, Imperial College, University of London, 1983. Lecturer, Department of Mechanical Engineering, University of Queensland, Australia. Applied & Numerical Mathematics [Rarefied Gas Dynamics and Kinetic Theory Methods for Continuum Gas Dynamics]. (January to June 1995)

Eric Morano - Ph.D., Engineering, University of Nice, France, 1992. Applied & Numerical Mathematics [Computational Aerodynamics]. (February 1994 to April 1995)

John R. Van Rosendale - Ph.D., Computer Science, University of Illinois, 1980. Applied & Numerical Mathematics [Floating Shock Fitted Euler Solution to Flow Past Airfoils]. (July 1994 to February 1995)

Nicolas G. Verhaagen - M.S., Aeronautics, Delft University of Technology, The Netherlands, 1970. Research Scientist, Faculty of Aerospace Engineering, Delft University of Technology, The Netherlands. Fluid Mechanics. (January 1995)

VII. SENIOR RESEARCH ASSOCIATE

David E. Keyes - Ph.D., Applied Mathematics, Harvard University, 1984. Computer Science. [Parallel Numerical Algorithms]

VIII. CONSULTANTS

Saul Abarbanel - Ph.D., Theoretical Aerodynamics, Massachusetts Institute of Technology, 1959. Professor, Department of Applied Mathematics, Tel-Aviv University, Israel. Applied & Numerical Mathematics [Numerical Analysis and Algorithms]

Ivo Babuska - Ph.D., Technical Science, Technical University, Prague, Czechoslovakia, 1951; Mathematics, Academy of Science, Prague, 1956; D.Sc., Mathematics, Academy of Science, Prague, 1960. Professor, Institute for Physical Science & Technology, University of Maryland. Applied & Numerical Mathematics [Finite Element Methods Associated With Structural Engineering]

Ponnampalam Balakumar - Ph.D., Aeronautics and Astronautics, Massachusetts Institute of Technology, 1986. Associate Professor, Department of Aerospace Engineering, Old Dominion University. Fluid Mechanics [Stability and Transition]

H. Thomas Banks - Ph.D., Applid Mathematics, Purdue University, 1967. Professor, Department of Mathematics, Center for Research in Scientific Computations, North Carolina State University. Applied & Numerical Mathematics [Control Theory]

Alvin Bayliss - Ph.D., Mathematics, New York University, 1975. Associate Professor, Technological Institute, Northwestern University. Fluid Mechanics [Acoustics]

John D. Buckmaster - Ph.D., Applied Mathematics, Cornell University, 1969. Professor, Department of Aeronautical and Astronautical Engineering, University of Illinois. Fluid Mechanics [Mathematical Combustion]

John A. Burns - Ph.D., Mathematics, University of Oklahoma, 1973. Professor, Virginia Polytechnic Institute and State University. Applied & Numerical Mathematics [Numerical Methods in Feedback Control and Parameter Estimation]

Gianfranco Ciardo - Ph.D., Computer Science, Duke University, 1989. Assistant professor, The College of William & Mary. Computer Science [Reliability Models]

Thomas C. Corke - Ph.D., Mechanical and Aerospace Engineering, Illinois Institute of Technology, 1981. Professor, Department of Mechanical and Aerospace Engineering, Armour College of Engineering, Illinois Institute of Technology. Fluid Mechanics [Instability Transition]

Ayodeji O. Demuren - Ph.D., Mechanical Engineering, Imperial College London, United Kingdom, 1979. Associate Professor, Department of Mechanical Engineeringand Mechanics, Old Dominion University. Fluid Mechanics [Numerical Modeling of Turbulent Flows]

Dennis B. Gannon - Ph.D., Mathematics, University of California, Davis, 1974. Professor, Department of Computer Science, Indiana University @ Bloomington. Computer Science [Investigation of Algorithms and Programming Techniques for Parallel Computers]

James F. Geer - Ph.D., Applied Mathematics, New York University, 1967. Professor, Systems Science and Mathematical Sciences, Watson School of Engineering, Applied Science and Technology, SUNY-Binghamton. Fluid Mechanics [Perturbation Methods and Asymptotic Expansions of Solutions to Partial Differential Equations]

David I. Gottlieb - Ph.D., Numerical Analysis, Tel-Aviv University, Israel, 1972. Professor, Division of Applied Mathematics, Brown University. Applied & Numerical Mathematics [Boundary Conditions for Hyperbolic Systems]

Chester E. Grosch - Ph.D., Physics - Fluid Dynamics, Stevens Institute of Technology, 1967. Professor, Department of Computer Science and Slover Professor, Department of Oceanography, Old Dominion University. Fluid Mechanics [Computational Fluid Mechanics and Algorithms for Array Processor Computers]

Max Gunzburger - Ph.D., Applied Mathematics, New York University, 1969. Professor, Department of Mathematics, Virginia Polytechnic Institute & State University. Applied & Numerical Mathematics [Fluid Dynamic Control]

Philip Hall - Ph.D., Mathematics, Imperial College, England, 1973. Professor, Department of Mathematics, University of Manchester, England. Fluid Mechanics [Instability and Transition]

Gene J.-W. Hou - Ph.D., Computational Mechanics, Design Optimization, University of Iowa, 1983. Associate Professor, Mechanical Engineering Department, Old Dominion University. Applied & Numerical Mathematics [Computational Mechanics Design Optimization]

Fang Q. Hu - Ph.D., Applied Mathematics, Florida State University, 1990. Assistant Professor, Department of Mathematics and Statistics, Old Dominion University. Fluid Mechanics [Computational Aeroacoustics] .

Harry F. Jordan - Ph.D., Physics, University of Illinois, 1967. Professor Department of Electrical and Computer Engineering, University of Colorado at Boulder. Computer Science [Parallel Computation]

Ashwani K. Kapila - Ph.D., Theoretical and Applied Mechanics, Cornell University, 1975. Associate Professor, Department of Mathematical Sciences, Rensselaer Polytechnic Institute. Fluid Mechanics [Mathematical Combustion]

Ken Kennedy - Ph.D., Computer Science, New York University, 1971. Chairman, Department of Computer Science, Rice University. Computer Science [Parallel Compilers and Languages]

Doron Kishoni - Ph.D., Theoretical and Applied Mechanics, Cornell University, 1983. Professor & Sr. Research Scientist, Physics Department & Applied Science Program, College of William & Mary. Fluid Mechanics [Computational Methods to wave propagation for Acoustics Electromagnetizm and Elasticity]

David A. Kopriva - Ph.D., Applied Mathematics, University of Arizona, 1982. Associate Professor, Department of Mathematics and SCRI, Florida State University. Applied & Numerical Mathematics [Development and Implementation of Spectral Methods for Viscous Compressible Flows]

Heinz-Otto Kreiss - Ph.D., Mathematics, Royal Institute of Technology, Sweden, 1960. Professor, Department of Applied Mathematics, California Institute of Technology, Applied & Numerical Mathematics [Numerical Analysis]

David G. Lasseigne - Ph.D., Applied Mathematics, Northwestern University, 1985. Assistant Professor, Department of Mathematics and Statistics, Old Dominion University. Fluid Mechanics [Computational Fluid Dynamics]

Anthony Leonard - Ph.D., Nuclear Engineering, Stanford University, 1963. Professor of Aeronautics, California Institute of Technology. Fluid Mechanics [Vortex Methods]

Scott T. Leutenegger - Ph.D., Computer Science, University of Wisconsin-Madison, 1990. Assistant Professor, Department of Mathematics and Computer Science, University of Denver. Computer Science [System Software Related to Databases for Scientific Data]

Robert W. MacCormack - M.S., Mathematics, Stanford University. Professor, Department of Aeronautics and Astronautics, Stanford University. Fluid Mechanics [Computational Fluid Dynamics and Numerical Analysis]

Kurt Maly - Ph.D., Computer Science, Courant Institute, New York University, 1973. Kaufman Professor and Chair, Department of Computer Science, Old Dominion University. Computer Science [High Performance Communication]

James E. Martin - Ph.D., Applied Mathematics, Brown University, 1991. Assistant Professor, Department of Mathematics, Christopher Newport University. Fluid Mechanics [Aeroacoustics]

David J. Middleton - Ph.D., Computer Science, University of North Carolina-Chapel Hill, 1986. Computer Science [Parallel Computing]

Sanjoy K. Mitter - Ph.D., Electrical Engineering, Imperial College of Science & Technology, London, 1965. Professor of Electrical Engineering, Co-Director, Laboratory for Information and Decision Systems, Director, Center for Intelligent Control Systems, Massachusetts Institute of Technology. Fluid Mechanics [Control Theory]

David M. Nicol - Ph.D., Computer Science, University of Virginia, 1985. Professor, Department of Computer Science, College of William and Mary. Computer Science [Mapping Algorithms onto Parallel Computing Systems]

R. A. Nicolaides - Ph.D., Computer Science, University of London, 1972. Professor, Department of Mathematics, Carnegie-Mellon University. Applied & Numerical Mathematics [Numerical Solution of Partial Differential Equations]

James M. Ortega - Ph.D., Mathematics, Stanford University, 1962. Professor and Chairman, Department of Applied Mathematics, University of Virginia. Computer Science [Numerical Analysis of Methods for Parallel Computers]

Demetrius Papageorgiou - Ph.D., Mathematics, University of London, 1985. Assistant Professor, Department of Mathematics, New Jersey Institute of Technology. Fluid Mechanics [Instability and Transition]

Anthony T. Patera - Ph.D., Applied Mathematics, Massachusetts Institute of Technology, 1982. Professor, Department of Mechanical Engineering, Massachusetts Institute of Technology. Fluid Mechanics [Surrogate Methods in Interdisciplinary Studies]

Ugo Piomelli - Ph.D., Mechanical Engineering, Stanford University 1987. Professor, Department of Mechanical Engineering, University of Maryland. Fluid Mechanics [Subgrid Scale Reynold's Stress Modelling and Large Eddy Simulation of Turbulent Flows]

Alex Pothen - Ph.D., Applied Mathematics, Cornell University, 1984. Professor, Department of Computer Science, Old Dominion University. Computer Science [Parallel Numerical Algorithms]

Daniel A. Reed - Ph.D., Computer Science, Purdue University, 1983. Assistant Professor, Department of Computer Science, University of Illinois. Computer Science [Parallel Processing]

Eli Reshotko - Ph.D., Aeronautics and Physics, California Institute of Technology, 1960. Kent H. Smith Professor of Engineering, Case Western Reserve University. Fluid Mechanics [Instability and Transition to Turbulence]

Philip L. Roe - Ph.D., Aeronautics, University of Cambridge, United Kingdom, 1962. Professor, Department of Aerospace Engineering, University of Michigan. Applied & Numerical Mathematics [Numerical Analysis and Algorithms]

Ahmed H. Sameh - Ph.D., Civil Engineering, University of Illinois, 1968. Head, William Norris Chair, and Professor, Department of Computer Science, University of Minnesota. Computer Science [Parallel Numerical Algorithms]

Sutanu Sarkar - Ph.D., Mechanical and Aerospace Engineering, Cornell University, 1988. Professor, Department of AMES, University of California, San Diego. Fluid Mechanics [Turbulence in High-Speed Compressible Flows]

Chi-Wang Shu - Ph.D., Mathematics, University of California, Los Angeles, 1986. Associate Professor, Division of Applied Mathematics, Brown University. Applied & Numerical Mathematics [Numerical Partial Differential Equations and Computational Fluid Dynamics]

Ralph C. Smith - Ph.D., Mathematics, Montana State University, 1990. Assistant Professor, Department of Mathematics, Iowa State University. Applied & Numerical Mathematics [Optimal Control Techniques for Structural Acoustics Problems]

Charles G. Speziale - Ph.D., Aerospace and Mechanical Sciences, Princeton University, 1978. Professor, Aerospace & Mechanical Engineering Department, Boston University. Fluid Mechanics [Turbulence Modeling]

Alan L. Sussman - Ph.D., Computer Science, Carnegie Mellon University, 1991. Assistant Professor, Department of Computer Science, University of Maryland. Computer Science [Parallel Computing]

Shlomo Ta'asan - Ph.D., Applied Mathematics, The Weizmann Institute of Science, 1985. Professor, Department of Mathematics, Carnegie Mellon University. Applied & Numerical Mathematics [Numerical Analysis and Algorithm Development]

Christopher Tam - Ph.D., Applied Mechanics, California Institute of Technology, 1966. Professor, Department of Applied and Numerical Mathematics, Florida State University. Fluid Mechanics [Computational Aeroacoustics]

Siva Thangam - Ph.D., Mechanical Engineering, Rutgers University, 1980. Professor, Department of Mechanical Engineering, Stevens Institute of Technology. Fluid Mechanics [Turbulence Modeling and Simulation]

Lu Ting - Ph.D., Aeronautics, New York University, 1951. Professor, Courant Institute of Mathematical Sciences, New York University. Fluid Mechanics [Acoustics]

Kishor Trivedi - Ph.D., Computer Science, Universtiy of Illinois-Urbana, 1974. Professor, Department of Electrical Engineering, Duke University. Computer Science [Performance and Reliability Modeling Methods, Tools and Applications]

Eli Turkel - Ph.D., Applied Mathematics, New York University, 1970. Associate Professor, Department of Applied Mathematics, Tel-Aviv University, Israel. Applied & Numerical Mathematics [Numerical Solution of Partial Differential Equations on Fluid Dynamics]

George M. Vahala - Ph.D., Physics, University of Iowa, 1972. Professor, Department of Physics, The College of William & Mary. Fluid Mechanics [Group Renormalization Methods for Turbulence Approximation]

Bram van Leer - Ph.D., Theoretical Astrophysics, Leiden State University, The Netherlands, 1970. Professor, Department of Aerospace Engineering, University of Michigan. Applied & Numerical Mathematics [Computational Fluid Dynamics]

Forman A. Williams - Ph.D., Engineering Sciences, California Institute of Technology, 1958. Professor, Department of Applied Mechanics and Engineering Sciences and Director, Center for Energy and Combustion Research, University of California, San Diego. Fluid Mechanics [Combustion]

Hans Zima - Ph.D., Mathematics, University of Vienna, Austria, 1964. Professor, Computer Science Department, University of Vienna, Austria. Computer Science [Compiler Development for Parallel and Distributed Multiprocessors]

IX. GRADUATE STUDENTS

David C. Cronk - Graduate Student at College of William and Mary. (August 1993 to Present)

Bryan K. Hess - Graduate Student at College of William and Mary. (March 1994 to Present)

Angelo Iollo - Graduate Student at Politecnico di Torino. (November to December 1994)

Frank P. Kozusko - Graduate Student at Old Dominion University. (October 1993 to Present)

Julia Levin - Graduate Student at Columbia University. (December 1994 to January 1995)

Joe L. Manthey - Graduate Student at Old Dominion University. (September 1993 to Present)

Patricia Slechta-Shah - Graduate Student at Boston University. (January 1994 to December 1994)

Willem Vermeer - Graduate Student at Delft University of Technology. (January 1995 to Present)

Robert V. Wilson - Graduate Student at Old Dominion University. (October 1992 to Present)

Kyle J. Winn - Graduate Student at College of William & Mary. (February 1994 to Present)

REPORT DOCUMENTATION PAGE

Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.

1. AGENCY USE ONLY *(Leave blank)*	2. REPORT DATE June 1995	3. REPORT TYPE AND DATES COVERED Contractor Report

4. TITLE AND SUBTITLE
Semiannual Report.
October 1, 1994 through March 31, 1995

5. FUNDING NUMBERS
C NAS1-19480
WU 505-90-52-01

6. AUTHOR(S)

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)
Institute for Computer Applications in Science
and Engineering
Mail Stop 132C, NASA Langley Research Center
Hampton, VA 23681-0001

8. PERFORMING ORGANIZATION REPORT NUMBER

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)
National Aeronautics and Space Administration
Langley Research Center
Hampton, VA 23681-0001

10. SPONSORING/MONITORING AGENCY REPORT NUMBER
NASA CR-198172

11. SUPPLEMENTARY NOTES
Langley Technical Monitor: Dennis M. Bushnell
Final Report

12a. DISTRIBUTION/AVAILABILITY STATEMENT

Unclassified–Unlimited

Subject Category 59

12b. DISTRIBUTION CODE

13. ABSTRACT *(Maximum 200 words)*
This report summarizes research conducted at the Institute for Computer Applications in Science and Engineering in applied mathematics, fluid mechanics, and computer science during the period October 1, 1994 through March 31, 1995.

14. SUBJECT TERMS
applied mathematics; numerical analysis; fluid mechanics; computer science

15. NUMBER OF PAGES
80

16. PRICE CODE
A05

17. SECURITY CLASSIFICATION OF REPORT Unclassified	18. SECURITY CLASSIFICATION OF THIS PAGE Unclassified	19. SECURITY CLASSIFICATION OF ABSTRACT	20. LIMITATION OF ABSTRACT